

THE LOTTERY PUZZLE AND PRITCHARD'S SAFETY ANALYSIS OF KNOWLEDGE

MARK McEVOY
HOFSTRA UNIVERSITY

ABSTRACT: The safety analysis of knowledge, due to Duncan Pritchard, has it that for all contingent propositions, p , S knows that p iff S believes that p , p is true, and (the "safety principle") in most nearby worlds in which S forms his belief in the same way as in the actual world, S believes that p only if p is true. Among the other virtues claimed by Pritchard for this view is its supposed ability to solve a version of the lottery puzzle. In this paper, I argue that the safety analysis of knowledge in fact fails to solve the lottery puzzle. I also argue that a revised version of the safety principle recently put forward by Pritchard fares no better.



In a number of recent publications Duncan Pritchard has outlined and defended what he calls the safety analysis of knowledge.¹ According to this view, for all contingent propositions, p , S knows that p iff S believes that p , p is true, and (the "safety principle"):

- (S1) In most nearby worlds in which S forms his belief in the same way as in the actual world, S believes that p only if p is true (Pritchard 2005, 156).²

This is an externalist analysis of knowledge, which can perhaps best be understood as a descendant of the tracking account of knowledge, proposed by Nozick and Dretske (though there are, as we shall see, significant differences between Pritchard's proposal and the tracking analyses of Nozick and Dretske).³ While I am in broad sympathy with Pritchard's externalist approach, I do not think that his analysis solves all of the problems he believes it to; specifically I find the safety analysis inadequate to a version of the lottery puzzle which Pritchard believes it can solve. In this paper, following exposition of the main points of, and motivation for, the safety analysis, I shall show that it does not solve this version of the puzzle.

I. SAFETY AND SENSITIVITY

Pritchard describes the Dretske-Nozick view as being committed to the sensitivity principle: if S knows that p, then S does not believe that p in the nearest possible world(s) in which not-p (Pritchard 2005, 48). This condition enables the sensitivity theorist to argue that we can know lots of the everyday propositions that we take ourselves to know (since in possible worlds where these obtain, we would believe them, and in possible worlds where they do not obtain, we would not believe them), and that we do not know that we are not brains in vats, since in the worlds where this is not true, we would believe that it is true.⁴

We can situate the safety analysis with respect to the sensitivity analysis by focusing on two of the main considerations that guide Pritchard's theorizing: the desire to exclude from knowledge, in a principled way, beliefs that are true as a matter of luck (Pritchard 2005, 1–2), and a commitment to what he labels the “core relative alternatives intuition”: “in order to know one only needs to be able to rule out the error-possibilities that are (in some sense . . .) relevant” (Pritchard 2005, 35). The former leads him to agree with the sensitivity theorist's verdict on everyday knowledge; the latter leads him to reject the sensitivity theorist's handling of denials of skeptical hypotheses. With respect to the former point, consider a couple of Pritchard's examples. First, an intuitive case of knowledge: I know, let us grant, that I am currently seated. The safety account yields this verdict, since in nearby possible worlds I am seated, and such worlds differ only in minor details, such as my looking out the window rather than looking at the screen (Pritchard 2005, 72). To say that my belief here is safe is, amongst other things, to say that it is not simply true as a matter of luck. In contrast to this, if my bank has just staved off a hacker's attempt to drain my funds, then although my belief that I have money in the bank is true, it is not knowledge. Again, the safety analysis gives the correct verdict: there are nearby possible worlds where my belief is not true, but where I would continue to believe it—worlds in which the hacker was one step ahead of the bank—and so my belief is not safe with respect to nearby worlds; it is only true as a matter of luck, and so, properly, it does not count as knowledge.

The safety analysis thus categorizes everyday beliefs similarly to the sensitivity analysis. However, Pritchard's interpretation of the core relevant alternatives intuition leads him to part company with the sensitivity theorist over the latter's treatment of denials of skeptical hypotheses. This core intuition limits the range of worlds that are relevant to the epistemic evaluation of everyday beliefs, such as the belief that I have hands. As we have seen, the sensitivity analysis holds that (presumably remote) worlds in which skeptical hypotheses obtain are not relevant to the epistemic evaluation of our *everyday* beliefs. However, Pritchard goes further, urging that the best interpretation of the core relevant alternatives intuition is that skeptical hypotheses, insofar as they concern “far-fetched error possibilities” are not relevant to knowledge *of any sort*—not even the knowledge “that one is not the victim of a skeptical scenario” (Pritchard 2005, 70 and 73). Our beliefs need only

hold true in *nearby* worlds.⁵ If our beliefs do so hold true, they are thus “safe” with respect to the truth conditions in nearby worlds, and so count as knowledge. The point, in a nutshell, is that if our world is not actually a skeptic-world, and if my belief that it is not is formed in an epistemically acceptable way (say through my reasonably reliable senses operating more or less as I believe them to), then such ex-hypothesi remote skeptic-worlds seem too remote to be at all relevant to the epistemic evaluation of my belief (Pritchard 2005, 70 and 75). The safety analysis, then, agrees with the sensitivity theorist that we know such things as that we have hands, but disagrees with the verdict that we do not know that we are not brains in vats.⁶ Pritchard (2005, 73) takes this to show that the safety analysis is truer to externalist intuitions than is the sensitivity analysis.

Pritchard (2005, 160) offers a further reason why skeptical worlds should not be taken into account in the epistemic evaluation of our beliefs in denials of skeptical hypotheses. Whether one is a sensitivity theorist or a safety theorist, once one has included in one’s analysis a requirement on relevant nearby worlds that a belief be formed *by the same method* as it is in the actual world, and assuming the actual world is not in fact a skeptical world, it follows that in skeptical worlds we would not form our belief that we are not in a skeptical world in the same way as we do in the actual world (in the skeptical world our belief is formed by a process of electro-chemical stimulation by vat-scientists). Since the method of belief formation differs from the skeptical world to the actual world, skeptical worlds fall outside the range of worlds relevant to the evaluation of our actual world beliefs.

This rejection of the requirement that our beliefs in the denials of skeptical hypotheses take account of skeptical worlds leads to one final difference between the sensitivity and safety views: the safety view maintains the closure of knowledge under known entailment. Closure is lost when one holds both (a) that one knows at least some everyday propositions and (b) that one fails to know the denials of skeptical hypotheses entailed by items included in (a). Since Pritchard’s safety view accepts (a) but rejects (b), the primary motivation for the denial of closure is undermined.

II. THE LOTTERY PUZZLE

There are many versions of the lottery puzzle; the version Pritchard discusses goes as follows. Suppose you buy a ticket in a fair lottery in which there are one million tickets. Your belief that your ticket will not win has only a one in a million chance of being false. From a purely probabilistic viewpoint, your belief that your ticket will lose stands a much greater chance of being correct than very many of our other (justified) beliefs in contingent truths. And yet, even if your belief turns out to be true, it has seemed to many—including Pritchard—that this belief is not knowledge. In contrast to this, according to Pritchard, when you read the result in a “normally reliable” newspaper which confirms that your ticket has lost, and form your belief that you have lost purely on the basis of the newspaper report? “this will (usually at least) suffice for knowledge” (Pritchard 2005, 162) that you have lost,

despite the fact that from a probabilistic point of view, there is a far greater chance of a newspaper error than there is of your initial belief (based on the lottery odds) of being mistaken.⁸ Hence the puzzle.

III. THE SAFETY SOLUTION TO THE LOTTERY PUZZLE

Pritchard notes that (S1) does not solve the lottery puzzle. This is due to the presence of nearby worlds in which the ticket-buyer forms the belief, based on the odds, that he will lose, but where this belief turns out to be false. The presence of such worlds implies that his actual world belief that he has lost is, if true, true in part as a matter of luck, and thus cannot be knowledge (Pritchard 2005, 163). To cope with this problem, Pritchard strengthens the safety principle as follows:

- (S2) For any contingent proposition, p , if S knows that p then, *in nearly all (if not all) nearby possible worlds in which S forms her belief about p in the same way as she forms her belief in the actual world, S only believes that p when p is true.* (Pritchard 2005, 163, emphasis added)

According to Pritchard, this solves the puzzle because:

The agent who forms her belief that she has lost the lottery purely on the basis of the odds involved lacks knowledge because her belief, whilst true and matching the truth in most nearby possible worlds in which she forms her belief in the same way as in the actual world, does not match the truth in a small cluster of nearby possible worlds in which what she believes is false (i.e., where she wins the lottery). (Pritchard 2005, 163)

Since her belief is not safe with respect to these nearby worlds, the truth of her belief involves an unacceptable element of luck and so that belief is not knowledge. Pritchard thus accounts for the common intuition that the belief that one will lose (based solely on the odds) is not, even if true, an instance of knowledge.

IV. WHY THE SOLUTION DOES NOT WORK

There are a number of worries one might raise for Pritchard's solution. One could, following Greco, argue that (S2) either fails to solve the lottery, or else rules out clear cases of knowledge.⁹ Greco's argument takes the following line: Suppose we take (S2) to require that our beliefs be safe only with regards to "nearly all" nearby possible worlds. In this case, our belief that we have lost the lottery (based solely on consideration of the odds) will count as knowledge. For, granted odds of, say, one million to one against, then my odds-based belief that I have lost will be false in only one of a million worlds. That is, this belief will be true in *nearly all* nearby possible worlds, and will thus count as knowledge, contrary to Pritchard's claim. Suppose, however, that to counter this problem, we require that our beliefs be safe with respect to all (as opposed to nearly all) nearby worlds. With the safety principle so interpreted, we get the desired result that my odds-based belief that I have lost the lottery is not an instance of knowledge (because

of the one nearby world in which I win). However, (S2) now excludes cases that intuitively seem to be instances of knowledge. Greco cites Sosa's "chute" case in this connection. In this case, S drops a bag down his building's garbage chute, and then forms the belief that the bag is in the basement garbage room. There is, however, an extremely unlikely possibility (which does not in fact obtain) that the garbage bag has instead snagged and is now stuck in the chute.¹⁰ Since there are nearby possible worlds in which S's belief is false, (S2), on this interpretation, seems to rule that S does not know that the bag is in the basement. Since this seems the wrong verdict, the argument concludes that when we interpret the safety principle in such a way as to enable it to solve the lottery puzzle, it fails with standard cases of knowledge.

Thus Greco. I want to put pressure in a different place, however. I want to argue that (S2) fails on its own terms as a solution to the lottery puzzle, regardless of its success or otherwise with other cases. Moreover, if my argument is sound, there is no need to exploit the ambiguity in (S2) between 'all' and 'nearly all,' since (S2) fails even when we stick to the 'nearly all' reading. The problem that I wish to raise is that while (S2) captures the intuition that my odds-based belief is not knowledge, due to the presence of nearby worlds in which that belief is false, (S2) also has the consequence that one does not obtain knowledge that one has lost even when one reads the result in a newspaper. Consider once more the reasoning behind the verdict that one's belief that one will lose, produced by reflection on the odds, does not count as knowledge: there is a small cluster of nearby worlds in which one forms this belief in the same manner, but where this belief turns out false. The belief in question fails to be true at nearly all nearby worlds and so is not safe. Exactly the same problem obtains with forming one's belief that one has not won by the method of reading the newspaper. There are, for example, worlds in which the newspaper contains a misprint. In most of these worlds, my belief that I have lost is true. However, in a small cluster of these worlds, I do win the lottery, and so my belief is false. Since my belief that I have lost is not true in nearly all nearby worlds, (S2) does not classify it as knowledge, contrary to Pritchard's claim.

There seem only three ways of rescuing Pritchard's solution to the puzzle. First, one could argue that the "misprint" worlds are sufficiently remote so that one's belief that one has lost is not required to be safe with respect to them (while simultaneously maintaining that worlds in which the corresponding odds-based belief is false are not so remote). Second, one could amend the safety principle so as to allow closer worlds to carry more weight than more distant worlds, and argue that the amended safety principle solves the lottery puzzle in an intuitively satisfying way.¹¹ Finally, one could argue that the method of forming one's belief (by reading the newspaper) differs from the actual world to the misprint worlds, and so one's belief that one has lost is not required to be safe with respect to those misprint worlds. In the remainder of this paper, I will argue that none of these responses can be maintained.

V. MODAL DISTANCE

Consider my actual world belief, formed by considering the odds, that I have lost. Let us call worlds in which this belief, formed by the same process, turns out to be false, but which are otherwise identical to the actual world (except for whatever changes are required to render this belief false) type (i) worlds. Type (i) worlds, one might urge, are sufficiently nearby that my actual world belief (based on the odds) that I have lost is required to be safe with respect to them. Consider now my actual world belief that I have lost, formed by reading the results in the newspaper. Type (ii) worlds are worlds in which I form this same belief, say through reading a misprinted newspaper, but where my belief is false. Except for this difference, along with whatever is required to cause this difference, type (ii) worlds are identical to the actual world. One might attempt to rescue (S2)'s solution to the lottery puzzle by suggesting that type (ii) worlds are sufficiently remote that my actual world belief that I have lost (formed by reading the newspaper) is not required to be safe with respect to them.

How, though, is this suggestion to be maintained? It might seem that there is an obvious way to do this: type (i) worlds differ from the actual world only in virtue of including one unlikely event, whereas type (ii) worlds differ from the actual world in virtue of including two unlikely events (my winning the lottery, and the newspaper including a misprint in the lottery results). This difference between these two types of worlds, it might be said, enables Pritchard to maintain that type (ii) worlds are sufficiently remote such that my belief that I have lost is not required to be safe with respect to them.

Pritchard, however, cannot use this argument as it stands. For he holds that an event's being unlikely is precisely *not* a criterion of it being modally distant. Pritchard explicitly states that while winning the lottery is a statistically rare event, it is also modally near: "The very attraction of a fair lottery lies in the fact that the possible world in which one wins is very like the actual world, even though it is in fact unlikely that such a possible world should be the actual world." (Pritchard 2005, 164; 2004, 197) Pritchard contrasts this rare-but-modally-near possibility with "a rare possibility . . . that is modally far off." Thus, the mere fact that type (ii) worlds contains one more unlikely event than type (i) worlds does not imply that type (ii) worlds are thereby more distant.

However, the argument can be altered so as to accommodate this aspect of Pritchard's thought. Pritchard explains the proximity of type (i) worlds on the basis of the fact that very little would need to change between these worlds and the actual world: "for example, a few numbered balls just need to fall into slightly different holes in the machine that draws the lottery numbers" (Pritchard 2004, 197). Elsewhere, in a discussion of a different case, Pritchard holds that worlds which are modally distant are those worlds in which "a lot would have to be different" (Pritchard 2005, 165). This gives us a second way of attempting to maintain the idea that type (i) worlds are modally nearer than type (ii) worlds: a lot would have to change between type (ii) worlds and the actual world, but very little would

need to change between type (i) worlds and the actual world; type (ii) worlds are therefore more distant than type (i) worlds. My belief that I have lost thus needs to be safe with respect to type (i) worlds, but not with respect to type (ii) worlds, and (S2)'s solution to the lottery can be reinstated.

But this argument too fails. For if the criterion for a world's nearness is simply that very little would need to change from it to the actual world, then surely many type (ii) worlds are *very* near. For, in addition to the few numbered balls falling into slightly different holes, all that is needed is (e.g.) that the person typing up the lottery results hits a key beside the one he should have hit. If this is correct, then type (ii) worlds are not sufficiently distant to be ignored, and the attempted rescue of (S2)'s solution to the lottery puzzle fails.

One might here object that much more needs to happen to produce the misprint than I am allowing. One might, for example, note that since we are dealing with a reliable newspaper, then in order for a misprint to occur, the (presumably) diligent character of the journalist and copyeditor would need to change in order for the misprint to occur. With so much changing, it seems not unreasonable to maintain that type (ii) worlds are more distant than lottery winning worlds after all, and thus the original suggestion could be maintained: type (ii) worlds are modally more distant than are type (i) worlds, and my belief that I have lost is not required to be safe with respect to the former.

But this response requires too much of a normally reliable newspaper. Take, as an example of a normally reliable newspaper, *The New York Times*. It contains misprints every day, some of which are detailed in the paper's "Corrections" column (and some of which make for very entertaining reading). Here are some examples taken from the month of October 2006:

Because of an editing error, an article on Monday about two Tibetan refugees shot to death by Chinese border security forces as they were trying to enter Nepal referred incorrectly to visits to Nepal by the Dalai Lama, the exiled Tibetan spiritual leader. He has visited once, in 1981—not "often." (October 19th)

An article on Tuesday about raids on the home of Representative Curt Weldon's daughter, as part of a federal corruption investigation, misstated the year Mr. Weldon was first elected to Congress, in Pennsylvania's Seventh District. It was 1986, not 1984. (October 19th)

Because of an editing error, an article in *Business Day* yesterday about a decision by a federal judge to void the conviction of Kenneth L. Lay, a former Enron executive, misstated the amount of money in one reference that prosecutors had sought from Mr. Lay. It was \$43.5 million, not billion. (October 19th)

Because of an editing error, a front-page article on Tuesday about the sentencing of Lynne F. Stewart, the radical defense lawyer, for aiding terrorism misstated the number of her grandchildren. She has 14, not 3. (October 20th)¹²

The point of these examples is that misprints and other errors can and do occur in normally reliable newspapers. The fact that these misprints occur despite the good character of actual journalists and copy editors shows that such mistakes can occur without requiring transworld changes to the characters of the relevant agents.¹³ Why think that the report of the winning lottery numbers would be especially immune to such misprints? The fallibility of the newspaper employees is all that is required for misprints to occur in any area—and since this is a constant between the actual world and all nearby worlds, very little has to change between the actual world and type (ii) worlds.

Granted only that the lottery results are not especially immune to the possibility of error, the conclusion is that type (ii) worlds are not, after all, sufficiently distant that my belief that I have lost is not required to be safe with respect to them. (S2) thus fails to solve the lottery puzzle.

VI. A NEW SUGGESTION: WEIGHTING WORLDS

I have just argued that Pritchard's safety solution to the lottery puzzle cannot be saved by taking type (ii) worlds to be too distant for consideration. However, in response to the kinds of considerations raised by Greco (see above), Pritchard has suggested (though not clearly endorsed) a new version of the safety principle, which leads to a correspondingly new attempted solution to the lottery. This new suggestion allows for a belief to be false in certain quite near worlds and still count as knowledge, provided that it is true in all *very* nearby worlds. Here is the amended safety principle:

- (S3) *S*'s belief is safe *iff* in most near-by possible worlds in which *S* continues to form her belief about the target proposition in the same way as in the actual world, and in all very close near-by possible worlds in which *S* continues to form her belief about the target proposition in the same way as in the actual world, the belief continues to be true. (Pritchard 2007a; 2007b)

(S3) allows for the following attempted solution to the lottery problem. Type (i) worlds in which I form my belief by considering the odds, but win, *are* very near the actual world in which I form my belief by this method, but lose. Thus, there are very nearby worlds in which my belief, formed by the same process, will turn out false, and so my odds-based belief that I will lose does not count as knowledge. This (it might be argued) contrasts with type (ii) worlds, which because of the additional difference of the misprint, are nearby the actual world, but not *very* nearby. If this is correct, then the falsity of my belief in type (ii) worlds does not prevent my belief that I have lost from being knowledge (as it would if such worlds were very nearby).¹⁴

In addition to enabling this solution to the lottery puzzle, (S3) allows for the following response to the chute case (Pritchard 2007b). First, recall that the possibility of the bag snagging is supposed to be “incredibly unlikely.” Pritchard takes

this to mean that the case is supposed to be one in which there is an imperfection in the chute, but one that is so slight that it would be unlikely that a bag would catch on it, and situated in such a place where hardly any bag would be likely to come near it. On such an understanding of the case, while there are close possible worlds in which the bag snags, there are no *very* close worlds are worlds in which the bag snags. (S3) thus yields the conclusion that S does indeed know that the bag is in the basement. (Pritchard (2007b) contrasts this with a reading of the case in which bags are frequently “almost” snagging on their way down the chute. On such a reading, he urges, it is no longer plausible to say that S knows that his bag is in the basement, and, since on this reading, there are very close worlds in which the bag snags, the safety principle delivers the correct verdict that S lacks knowledge.)

There are two problems to note here. First, there is an issue here about how to generalize our judgments of modal distance in a systematic way. This is not yet a very significant problem, since in the case of both the lottery puzzle and the chute case, there is, it might be argued, intuitive application of the idea that worlds which involve more change from the actual world are more distant than worlds which involve less change. However, it should be noted that Pritchard has offered no systemic method to underwrite these judgments; this will become a serious problem below.

The second problem I want to raise here is an extension of a case discussed by Pritchard, in which one comes to know the time on the basis of looking at a (functioning) clock, in good cognitive conditions; Pritchard notes that under such conditions “in most nearby possible worlds where I form this same belief on the same basis, my belief will be true . . .” (Pritchard 2004, 210). Pritchard goes on to note that “it is certainly true that there will be a small cluster of possible worlds where the belief I form on this basis is false (such as where the clock stopped a few minutes before I looked)” (Pritchard 2004, 210). Nonetheless, he holds that despite the existence of such nearby worlds, we can obtain genuine knowledge from looking at the clock. This is in contrast to the world in which the clock we look at has just stopped, but luckily tells the right time. In this case, there are too many nearby worlds in which this method results in a false belief.

A variation on the clock case shows that (S3) is unacceptable. There are very near worlds in which the reliable, regular clock we look at would yield a true belief that is not safe. These include, for example, worlds in which the battery was inserted in the clock a couple of days earlier than it was in the actual world, worlds in which the battery had left the factory filled with slightly less power than in the actual world, or (to switch to electric clocks) worlds in which the fuse in the clock’s plug was a little older than in the actual world. In some such worlds, the clock had stopped just before I looked at it, because the battery had run down, or because the fuse in the clock plug blew out. To see how such worlds create a problem for (S3), consider how little changes between such worlds and the actual world. All that need change is that the battery powering the clock was inserted a day or two

earlier than in the actual world (or contained *slightly* less power from the outset than the actual world battery, or, in the case of the electric clock, that the fuse was *slightly* older when it was placed in the plug than was the case in the actual world). Worlds that differ from the actual world only with respect to these details are surely very near to the actual world. Due to the presence of such worlds, together with the fact that (S3) requires that for a belief to be knowledge, it must be true in all very near worlds, (S3) implies that we never obtain knowledge of the time by looking at a normally reliable clock. (S3) thus sets the standard for knowledge too high, and must, for this reason, be rejected.

One might, however, attempt to defend (S3) as follows. Given that Pritchard's strategy in dealing with the chute case was to argue that the worlds in which the bag gets snagged are not really *very* near, we should expect that he would try to argue that worlds in which the battery runs out, or in which the fuse dies, are likewise not very nearby worlds. This move had some plausibility with the chute case, since for the case to work, the imperfection in the chute had to be situated in such a way that it was all but impossible for bags to snag, and that because of this, worlds in which it did snag would have to be quite different from the actual world. However, the move has no plausibility whatever here. For surely worlds in which a fuse starts life a day or two before it does in the actual world, or worlds in which a battery leaves the factory with *slightly* less juice than it does in the actual world really are very near to the actual world. If one argues otherwise, it begins to look very much as though one is making *extremely* fine-grained distinctions with respect to modal distance. And at this point the lack of any systematic method of so doing becomes a *serious* deficiency, opening one up to the charge of ad hoc reasoning. In other words, without providing a systematic method of reckoning modal distance, Pritchard cannot argue that the problematic nearby worlds in this example are insufficiently nearby to count. Since Pritchard has no such method, (S3), along with its putative solution to the lottery puzzle, must be rejected.

VII. DOES THE PRESENCE OF ERROR ALTER THE METHOD?

One might attempt to rescue (either version of) the safety solution to the lottery puzzle by excluding type (ii) worlds from consideration on the basis that the method by which one arrives at one's belief that one has lost is different from the method by which it is arrived at in the actual world. After all, one could describe the method used in misprint worlds as "reading an erroneous newspaper report," and this would differ from the actual world method (recall that the newspaper in question is stipulated to be normally reliable). Given, then, that (each version of) the safety analysis of knowledge requires that a belief that *p* be safe only with respect to worlds in which the belief that *p* is formed *by the same method* as in the actual world, this proposal would exclude type (ii) worlds from the range of worlds relevant to the epistemic evaluation of one's belief that one has not won the lottery. This, in turn, could reinstate the safety solution to the lottery puzzle.

Whatever the merits of this proposal, it is not open to Pritchard. For Pritchard assumes elsewhere that the same type of method can involve an error in some, but not all, of its tokenings. In his discussion of the clock case, mentioned above, Pritchard notes that we can come to know the time on the basis of looking at a (functioning) clock, in good cognitive conditions. Pritchard notes that under such conditions “in most nearby possible worlds where I form this same belief on the same basis, my belief will be true . . .” (Pritchard 2004, 210). However, Pritchard goes on to note that “it is certainly true that there will be a small cluster of possible worlds where the belief I form *on this basis* is false (such as where the clock stopped a few minutes before I looked)” (Pritchard 2004, 210, emphasis added). Clearly we have here a case in which the belief-forming process is held to be identical across its tokenings regardless of whether there is sometimes an error involved in the input to the believer’s cognitive system. Nor is this the only such example. In a discussion of Goldman’s barn façades case,¹⁵ Pritchard, having first discussed the case where Henry forms his belief on the basis of looking at a (genuine) barn, goes on to observe that “in a great number of nearby possible worlds where he forms his belief *on the same basis*, Henry will be looking at barn facades rather than barns” (Pritchard 2004, 207, emphasis added). As both of these cases show, Pritchard cannot consistently rescue his solution to the lottery puzzle by denying that the method whereby one forms one’s belief by reading a mistaken newspaper is identical to the method used in the actual world.

Moreover, since the presence of error in type (ii) worlds is *ex hypothesi* the only difference between such worlds and the actual worlds that could at all plausibly be said to pertain to the method of belief production,¹⁶ if this difference does not produce a difference in methods of belief production between type (ii) worlds and the actual world, then the methods are identical. There is thus no possibility of avoiding the problems raised here by distinguishing methods of belief production.

VIII. CONCLUSION

Pritchard has suggested two versions of the safety principle—(S2) and (S3)—which, he holds, solve the lottery puzzle. (S2) explains the common intuition that my belief (based solely on the odds) that I will not win the lottery is not knowledge. However, since it achieves this because of the presence of very close worlds in which this belief turns out false, and since there are also nearby type (ii) worlds in which my belief that I have not won (based on reading the newspaper) is false, (S2) fails to solve the lottery, since it fails to count as knowledge my newspaper produced belief that I have lost. (S3), on the other hand, might seem to solve the puzzle, but it either fails to categorize intuitive cases of knowledge correctly, or else, if it attempts to make fine-grained distinctions in modal distance without providing a systematic method for so doing, is open to the charge of ad hoc reasoning.

Short of further amending the safety principle, there seem only two ways to rescue the safety solution to the lottery puzzle: one could argue either that type (ii) worlds, unlike type (i) worlds, are too remote and so can be disregarded, or

one could argue that the actual world method of reading the newspaper, and so coming to believe that I have lost, differs from the corresponding method used in type (ii) worlds. I have offered reason to think that neither of these responses has any hope of success. The conclusion, then, is that the safety analysis, whatever its other merits, fails to solve the lottery puzzle.¹⁷

ENDNOTES

1. See, for example, "Virtue Epistemology and Epistemic Luck," *Metaphilosophy* 34 (2003): 106–130; "Epistemic Luck," *Journal of Philosophical Research* 29 (2004): 129–222; *Epistemic Luck* (Oxford: Oxford University Press, 2005).
2. This is subject to a couple of complications, which I will discuss in connection with the lottery puzzle.
3. See, for example, Fred Dretske, "Epistemic Operators," *Journal of Philosophy* 67 (1970): 1007–1023; Fred Dretske, "Conclusive Reasons," *Australasian Journal of Philosophy* 49 (1971): 1–22; Robert Nozick, *Philosophical Explanations* (Cambridge, Mass: Belknap, Harvard University Press, 1981). The view is referred to in the literature as the tracking analysis of knowledge. However, as Pritchard's view shares many of the features of the tracking analysis, I will follow him in distinguishing between the Dretske-Nozick view (sensitivity) and his own view (safety).
4. Famously, this results in the denial of closure: I know that I have a hand, I know that my having a hand entails that I am not a brain in a vat, but yet I do not know that I am not a brain in a vat.
5. This, of course, assumes that there are no skeptical worlds near to the actual world. As Pritchard (2005, 72) notes, the safety analysis allows our everyday knowledge only if skeptical worlds are in fact remote (though as an externalist, he does not require that we be in a position to (reflectively) know this).
6. Though of course, we cannot have reflective knowledge, of the kind required by the epistemological internalist. Pritchard (2005, 77) is quite explicit about this. Since I will not be taking Pritchard to task on the internalist/externalist divide, nothing I say here will depend on this issue.
7. Pritchard (2005, 162) is quite explicit that the newspaper-produced belief be formed *purely* on the basis of the newspaper report, and I shall follow him in this. One might instead urge that the belief that one forms upon reading the newspaper is one in which consideration of the odds against one's winning also plays a part. However, if one reads the case in this way, the lottery puzzle disappears. For now we are comparing (i) the method of forming one's belief that one will lose purely through consideration of the odds with (ii) the method of forming this same belief content through (i) *in conjunction with* reading the confirming report in a usually reliable newspaper. As (ii) is clearly a more reliable method than (i), the conclusion that (ii) leads to knowledge, while (i) does not, does not have any air of paradox about it.
8. Pritchard doesn't specify what he has in mind by the phrase "usually at least." However, if it is to be consistent with the safety account of knowledge, it must allow that coming to believe that p by reading a report in a normally reliable newspaper will yield knowledge if

(a) the report is accurate and—subject to a further revision to be detailed below—(b) in most nearby worlds in which one arrives at this belief by this method (i.e., reading the report), one only believes p when p is true. I add this remark lest it be thought that this vagueness in the account could be pressed into service to respond to the objections that follow. As will become clear, it cannot.

9. John Greco, “Worries About Pritchard’s Safety,” *Synthese* 158 (2007): 299–302.

10. Greco discusses this case both in his 2007 and in his “Virtue and Luck, Epistemic and Otherwise,” *Metaphilosophy* 34.3 (2003): 361 (The case is originally due to Ernest Sosa. See his “Skepticism and Contextualism,” *Philosophical Issues* 10 [2000]: 13). Greco also discusses several other counterexamples to the safety principle, all with the same structure as the chute case: an agent forms a true belief regarding the failure to obtain of an unlikely possibility; the agent intuitively has knowledge, but as there are nearby worlds in which this belief is false, the safety principle seems forced to rule that the agent lacks knowledge.

11. Pritchard describes (without clearly endorsing) a version of this method in a couple of recent papers. See his “Anti-Luck Epistemology,” *Synthese* 158 (2007a): 277–297. See also “Knowledge, Luck, and Lotteries,” in *New Waves in Epistemology*, ed. V. F. Hendricks and D. H. Pritchard (London: Palgrave Macmillan, 2007b).

12. Here are a couple of other examples, both from October 20th: “An article in Business Day on Wednesday about Sony’s plans for its new game console, PlayStation 3, misspelled the surname of a co-chief operating officer of Sony Computer Entertainment America. He is Jack Tretton, not Trenton.” “A sports article in some copies yesterday about the Mets’ 4–2 victory over the St. Louis Cardinals in Game 6 of the National League Championship Series misspelled the surname of a Mets star from their 1969 championship team. He is Ron Swoboda, not Svoboda.” Grammatical errors too, though less frequent, are not unheard of. Here is one from the print edition of October 3rd, 2006: “For example, in Las Vegas, which is surrounded by desert, residents can get rebates for using covers on swimming pools to slow evaporation. requires less water” [sic].

13. Of course, one could stipulate that nothing counts as a “normally reliable newspaper” if it includes misprints, but this would merely be to define the problem away. For if this is what is meant by ‘normally reliable,’ then there is no such thing as a normally reliable newspaper, and a fortiori, no such thing as coming to discover that one has lost by reading the results in a normally reliable newspaper.

14. I am not convinced by this suggestion. It seems to me that type (ii) worlds are very nearby the actual world, and thus that (S3) fails for the same reason that (S2) fails. However, absent a systematic method of judging modal distance, I see no way to argue conclusively either for or against the close proximity of type (ii) worlds, and so will pursue a different line of argument here.

15. For the original presentation of this example, see Alvin Goldman, “Discrimination and Perceptual Knowledge,” reprinted in *The Theory of Knowledge*, ed. L. Pojman (Calif.: Wadsworth, 1999).

16. The other differences between type (ii) worlds and the actual world all have to do with the production of the misprint, rather than the cognitive process leading from perceptual input to belief formation. They thus cannot be part of the method of belief formation, and, a fortiori, cannot underwrite a difference between methods of belief formation from type (ii) worlds to the actual world.

17. I wish to thank my colleagues at the Philosophy Department at Hofstra for their helpful comments, made in response to an earlier version of this paper, which was presented at the departmental colloquia series. Those comments ensured that it was not the earlier version of the paper that was submitted. Thanks are also due to two anonymous referees for *The Journal of Philosophical Research* whose comments occasioned some important revisions.