CHAPTER 5:

LOTTERIES, INSENSITIVITY, AND CLOSURE

Contents

CHAPTER 5:

LOTTERIES, INSENSITIVITY, AND CLOSURE

1. The Harman Lottery Puzzle

In some lottery situations, the probability that your ticket is a loser can get very close to 1. Suppose, for instance, that yours is one of 20 million tickets, only one of which is a winner. Still, it seems that (1) You don't know yours is a loser and (2) You are in no position to flat-out assert that your ticket is a loser. 'It's probably a loser,' 'It's all but certain that it's a loser', or even, 'It's quite certain that it's a loser' seem alright to say, but, it seems, you are in no position to declare simply, 'It's a loser.' (1) and (2) are closely related phenomena. In fact, I'll take it as a working hypothesis that the reason 'It's a loser' is unassertable is that (a) you don't seem to know that your ticket is a loser, and (b) in flat-out asserting some proposition, you represent yourself as knowing it. (b), of course, is the knowledge account of assertion, in one of its forms, explained and defended (with references to other works in which it is further explained and defended) in Chapter 3 of Volume 1 (DeRose 2009: esp. 92-98). This working hypothesis will enable me to address these two phenomena together, moving back and forth freely between them, and interacting with the work of others, some of which has been addressed to each of these, though my main focus will be on (1).

Following Gilbert Harman (though I tweak his example a bit), we note that things are quite different when you report the results of last night's basketball game. Suppose your only source is your morning newspaper, which did not carry a story about the game, but simply listed the score, 'Knicks 83, at Bulls 95,' under 'Yesterday's Results.'[1] Now, it doesn't happen very frequently, but, as we all should

[1] Harman contrasts our apparent lack of knowledge in the lottery case with the knowledge we often apparently gain either by being told some fact or reading it in the newspaper. The use of a sports score here

suspect, newspapers do misreport scores from time to time. On a few occasions, my paper transposed a result, attributing to each team the score of its opponent. In fact, that your paper's got the present result wrong seems quite a bit more probable than that you've won the lottery of the above paragraph. Still, when asked, 'Did the Bulls win yesterday?,' 'Probably' and 'In all likelihood' seem quite unnecessary. 'Yes, they did,' seems just fine. The newspaper, fallible though it is, seems to provide you with knowledge of the fact that the Bulls won. Indeed, if you're asked whether you know if the Bulls won, you'll likely respond positively. And, still following Harman, we note that this combination of natural evaluations is quite puzzling. In a very revealing passage, where *N* is the number of tickets in the lottery, and the 'testimony case' is one where 'a person comes to know something when he is told about it by an eyewitness or when he reads about it in the newspaper,' Harman writes:

> A person can know in the testimony case but not in the lottery case, or so we would ordinarily and naturally judge. In the lottery case a person cannot know he will lose no matter how probable this is. The contrast between the two cases may seem paradoxical, since witnesses are sometimes mistaken and newspapers often print things that are false. For some *N*, the likelihood that a person will lose the lottery is higher than the likelihood that the witness has told the truth or that the newspaper is right. Our ordinary, natural judgments thus seem almost contradictory. How could a person know in the testimony case but not in the lottery case? (Harman 1968: 166)

Here Harman issues an insightfully strong statement of the apparent tension between these two individual judgments, but, curiously, Harman's own approach to the epistemological puzzle we face here seems to neglect that insight. What immediately follows the above passage is this:

> At this point many philosophers would reject one of the ordinary judgments no matter how natural the judgment may be. But such a rejection would be premature. My strategy is to ask how beliefs are based on reasoning in the two cases …. (1968: 166)

Harman takes it as a working assumption (that's not his own description, but I believe it accurately describes his way of proceeding) that the two 'natural'

---

is my own embellishment. These days, in contrast to when I wrote the paper that is the basis of the first half of this chapter, very few people I know get their sports results from newspapers, relying instead on the internet. I have decided to be out-of-date and stick with newspapers, however, not only to retain the connection with Harman's discussion, but also because printed, unchanging newspapers seem in some ways better analogues of lottery tickets than are web pages.

‘ordinary judgments’ here are correct, to see where that working assumption leads, and if it leads to a sensible enough picture of what’s going on, to take that as reason to accept the picture that emerges. And, as things turn out, Harman thinks this procedure does lead to a sensible destination, where we learn what kinds of grounds underwrite knowledge of flat-out claims (as opposed to claims about what is probably the case) of the likes that the Bulls won or that one has lost the lottery.[2]

Harman’s basic approach has been followed in most of the subsequent literature on the epistemology of lotteries, with most epistemologists working on the topic pretty much taking it for granted, or at least taking it as a working assumption, that subjects don’t know that they’ve lost the lottery (in the standard situation), while they do know various ordinary things (and so the denial of knowledge in lotteries isn’t part of some general skepticism).[3]

---

[2] Dana Nelkin reads me as proceeding similarly in (DeRose 1996), citing me as her example of the many who take denying that we know we’ve lost in the lottery to be the ‘obvious (or almost obvious)’ way of handling the lottery puzzle (Nelkin 2000: 375). But in that paper, which forms the basis of the first half of this chapter, I was just trying to account for the intuitions about the particular cases, explicitly leaving aside questions about what we really do and don’t know here (DeRose 1996: 568). As we’ll see in this chapter, especially as I get to my own solution in section 14, but also as I continue in this opening section, I’m very far indeed from taking a flat denial of knowledge as the obvious way to go.

[3] Rachel McKinnon reports that ‘A growing consensus has formed that [Ticket n will lose] and propositions like it are neither knowable nor assertible’ (McKinnon 2013: 524). I don’t know if this is quite a ‘consensus’, nor if it is growing. (In neither case do I mean to be suggesting things are otherwise. I really just don’t know.) However, it certainly does seem true that at least most working on the problem go that way, and are more interested in explaining how or why our intuitions about the individual cases are right (how we have knowledge about many ordinary things, but don’t know that we’ve lost the lottery) than in seriously investigating the thought that we might know that we’re losers of lotteries (in the standard situation). (And I apparently, and understandably, seemed to Nelkin to be among that majority in (DeRose 1996); see the previous note.)

McKinnon’s own interesting explanation for the contrast, made from within the Relevant Alternatives theory of knowledge, is based crucially on her claims that ‘one may properly ignore destabilizing alternative possibilities in coming to know or assert some proposition’ (2013: 538), which is key to how we know ordinary, non-lottery propositions that we’re fallible with respect to, while we cannot properly ignore alternatives that are not destabilizing, and her claim that one’s winning the lottery, while it would cause many changes in one’s beliefs, would not be, in the way McKinnon explains, *destabilizing* to one’s view of how the world works. See (2013: 537-43) for McKinnon’s account of destabilization. Following the procedure I will use through much of the first half of this chapter, my main worry about McKinnon’s particular proposal is that it would seem to give the wrong result when we modify the lottery situation in one of the most natural ways to do so in order to test her solution. Imagine, then, a subject who is otherwise in the standard lottery situation, but does have background beliefs that render ‘I’ve won the lottery’ destabilizing for her. (Perhaps it

But there would seem to be *three* claims concerning knowledge in these cases that are each individually very plausible: the two individual judgments that Harman is taking as his starting points, plus the comparative appraisal that those two individual judgments are at odds with one another—that if one does not know in the lottery case, then one does not know in the newspaper case. That last, comparative claim would also be a very natural and plausible one, as it would seem Harman would have to agree. After all, his own reaction to the initial attempt to put the two individual judgments together was to report that their combination seems not only 'paradoxical,' but 'almost contradictory'! And I don't think it's sensible to suppose that the initial intuitive plausibility of the individual judgments overwhelms that of the comparative judgment—not because the comparative judgment is so unshakably solid, but because those individual judgments turn out to be rather flighty themselves, as we'll discuss in sections 16-18. That flightiness, together with the considerable intuitive power of the comparative judgment, argues for viewing our puzzle as consisting of three claims, all of which we should seek to do justice to. Of course, once that third claim is added, we can no longer even hope that the manner in which we do justice to all the relevant claims will be to endorse all of them in any strong way, since this third claim says that the first two claims can't both be right. We will eventually have to engage in some kind of intuitive damage control. But that's to be expected when venturing near paradox.

Those who may need some help in appreciating the power behind the comparative judgment that I'm suggesting we add to the mix may do well to briefly consider the trouble a treatment like Harman's suffers. Let's look at the violence his account does to intuitive ties between knowledge and rational action. Suppose you are faced with a choice between two tickets, which are each tickets to (different) one million dollar lotteries. Unfortunately, they are each at least probably losers. With respect to what we can call the 'normal' ticket, you're in the situation usual to philosophical discussions of lotteries: You have the lottery-like, statistical types of grounds for thinking it's at least probably a loser that Harman thinks cannot produce knowledge of its being a loser. So, Harman will endorse the intuitive judgment that you don't know it's a loser. Things are different with what we can call

---

is central to her view of how the world works that God is watching over things, and would never allow her to win the lottery under the current circumstances.) Intuitively, this subject seems not to know that she's lost the lottery every bit much as does a normal subject who does not have such background beliefs, and so to whom 'I've won' is not destabilizing, in McKinnon's sense.

the 'weird' ticket; as could have been guessed from its label, its history is a bit more colorful. With respect to it, you have what Harman would rule are knowledge-producing grounds—let's say through some kind of testimony—that it is a loser. Suppose the weird ticket comes from the wallet of someone who recently died. This person was carrying in his wallet two lottery tickets, one of which was a winning ticket to a 1 million dollar lottery that he was planning to cash in soon, and the other of which was an old, confirmed loser of a ticket to some long-ago lottery that he was holding on to for sentimental reasons. The circumstances under which you have come to be offered a chance to take this weird ticket would have been such that you would have thought it had a 50/50 chance of being the winner—except that, sadly, you now have somebody's testimony—say, a niece of the owner of the wallet—that it's the old loser ticket. Adjust the nature and circumstances of this testimony so that in this situation it is good enough, but just barely good enough, to be knowledge-producing by Harman's account, and that as an agent you are aware of the features that render the testimony just barely knowledge-producing. This should result in one of those situations where, on Harman's account, what you believe by means of testimony—that the weird ticket is a loser—has a greater chance of being wrong than that you will lose a lottery in the standard lottery situation you're in with respect to the normal ticket, yet you know what you learn through testimony while you don't know that the ticket in the standard lottery situation will lose. We needn't wrangle over whether, nor over the details of how, this could come about. Harman admits that sometimes in these situations what you know is less probable from your point of view than that you've lost a lottery in a standard lottery situation, which you don't know, and he notes that this is puzzling, or worse. We are just imagining such a baffling (seemingly 'almost contradictory', by Harman's own lights) situation in just a little bit of detail. So, from your perspective, the weird ticket has a better chance of being a winner than does the normal ticket. Presumably, then, given the choice between them, you should choose the weird ticket. But Harman rules that you know that the weird ticket will not win, but you don't know that the normal ticket will not win. So, if you follow his account on the epistemology of lotteries, you will end up saying that in such a situation, where you are choosing between two tickets specifically toward the hope of winning 1 million dollars, and you know of one of them but not of the other that it will lose, you should choose the ticket *that you know will lose* over the one you do not know to be a loser! Well, either that or you will instead say that in this situation you actually should choose the ticket that, from your own point of view, has less of a chance of winning. Either way, the intuitive costs seem quite substantial indeed.

Note, however, that it isn't any peculiarity of Harman's account that generates this problem. Indeed, we didn't even get very far into Harman's own account of just what separates the lottery from the newspaper case. What generates Harman's problem is really just that he endorses the particular intuitive judgments concerning when we do and do not know in the individual cases involved, resulting in his taking a counter-intuitive stand against the third, comparative judgment that is also intuitively quite plausible here. Any account that so confirms the two intuitions that Harman validates will bear similar intuitive costs with respect to the third intuition in play here.

In the first portion of this chapter, through section 13, we will be focused just on explaining the two individual judgments constitutive of our puzzle: Why do we judge that we do have knowledge—and assertability—in cases like the newspaper case, but that we don't in cases like the lottery? Just when does our tendency to judge subjects to be ignorant kick in for lottery-like reasons? Here we'll be covering much the same terrain that Harman does, with the difference that we'll be more open to our investigation ending up at an explanation for our individual judgments that will eventually lead us to a conclusion on which we don't endorse both of those judgments. So we will at this opening stage sidestep questions about whether you really do know in our two cases by focusing on explaining why it (at least) *seems* to us that you know the relevant proposition in the newspaper, but not in the lottery, case. In the second portion of this chapter, starting with section 14, we will turn to the questions of how best to account for the intuitive data and we will there address the question of whether we really do know we've lost in various lottery situations.

But a point of clarification is needed before we get into our initial explanatory task: If, in the newspaper case, one were confronted by a skeptic determined to make heavy weather of the possibility that the paper has made a mistake, then one might be led to take back one's claim to know the Bulls have won, and to refrain from flat-out asserting that they won. Indeed, such a skeptic may prompt you to feel, not just a generic skeptical attitude toward your belief that the Bulls won, but a skepticism that has a distinctively lottery-like feel to it. She may stress that newspapers of course have misreported scores, and focus attention on how you could possibly know that this isn't one of those occasions. 'It's like I'm in a lottery!' (though a friendly one, in which most 'tickets' win), you might be moved to exclaim in dismay before admitting that, no, you guess you don't really know that the Bulls won. On the other hand, as we'll discuss in sections 16-18, there are rather anti-skeptical situations in which folks do seem moved to judge, sometimes quite seriously, that they do indeed know that they have lost the lottery. The data we're trying to explain is indeed rather flighty and dodgy in various ways. And we will want to get a handle on just why that is. But, for now, what we want to explain is why, *with no such a skeptic in sight*, we typically do judge that we know in the newspaper, but not in the lottery, case, and in other cases much like them. (Unless so judging in the lottery case makes *us* skeptics, in which case we want to know why we're so naturally skeptics in the lottery, but not in the newspaper, case.)

2. The Explanation: SCA

Although several candidate explanations for why we seem to lack knowledge in the lottery case (while possessing it in the newspaper case) suggest themselves quite naturally, I accept the subjunctive conditionals account (SCA) of this phenomenon—an account that may not immediately jump to mind. Indeed, one of my main goals in this chapter, in addition to my independent interest in solving Harman's lottery puzzle, is to further support SCA, which I also employ in SSP, by appeal to its ability to solve this lottery puzzle. According to SCA, the reason we judge that you don't know you've lost the lottery is that (a) although you believe you are a loser, we realize that you would believe this even if it were false (even if you were the winner), and (b) we tend to judge that S doesn't know that P when we think that S would believe that P even if P were false. By contrast, in the newspaper case, we do not judge that you would believe that the Bulls had won even if that were false (i.e., even if they hadn't won).

SCA is close to the explanation that Fred Dretske attempts at (Dretske 1971: 3-4) and is the explanation that would be suggested by Robert Nozick's theory of knowledge in (Nozick 1981). But one need not buy into Dretske's or Nozick's analyses of knowledge to accept SCA. As I stressed back in SSP, (b) is far from a set of necessary and sufficient conditions for knowledge; it posits only a certain block which prevents us *from judging* that subjects know. This is important because Dretske's and Nozick's analyses of knowledge imply strongly counter-intuitive failures of the principle that knowledge is closed under known entailment. The correctness of SCA has been obscured by its being tied to theories of knowledge with such unpleasant implications, and also because not much of an argument has been given in its favor. I hope to remedy this situation by applying SCA to a variety of lottery-like and newspaper- like cases in the first stage of our inquiry and arguing that it outperforms its rivals in terms of explaining our judgments about what is and isn't known. If I succeed in showing that SCA is the best explanation for why we have the particular intuitions we have, that should motivate us to seek an account of knowledge that makes sense of SCA without doing the violence to various comparative judgments we're inclined to make that Dretske's and Nozick's analyses do. In the second stage of this chapter, I'll present a solution to the puzzle, analogous to our solution of the skeptical problem in SSP, that does just that.

One reason to accept SCA is that other initially plausible accounts, including the ones that naturally come to mind, don't work, as I'll try to show in what follows. In the meantime, what is there to recommend SCA, other than the fact that it yields the desired distinction between our two cases?

First, there is (b)'s initial plausibility. Though SCA may not naturally jump to mind, once it is suggested, it seems to provide a very intuitive explanation. If it can be shown to us that a subject would believe something even if it were false, that intuitively seems a pretty strong ground for judging that the subject doesn't know the thing in question.

Second, as I noted in SSP, there is the consideration that, in the lottery situation, even the most minute chances of error seem to rob us of knowledge and of assertability. In light of this, it seems puzzling that we will judge that a subject does know she's lost the lottery after she's heard the winning numbers announced on the radio and has compared them with the sadly different numbers on her ticket. For the announcement could be in error; she might still be the winner. Unlikely, to be sure. But if even the most minute chances of error count, why does it seem to us that she knows now that the announcement's been heard? SCA's answer is that once our subject has heard the announcement, (a) no longer holds. We no longer judge that if our subject were the winner, she'd still believe she was a loser; rather, we judge that if she were the winner, she'd now believe that she was, or would at least be suspending judgment as she tried to double-check the match. The very occurrence which makes us change our judgment regarding whether our subject knows, no longer denying that she knows, also removes the block which SCA posits to our judging that she knows. This provides some reason for thinking that SCA has correctly identified the block.

Tied up with the above recommendations is the fact that SCA nicely explains a lot of puzzling intuitions to the effect that subjects don't know propositions to be the case, in examples not involving lotteries and, as I stressed back in SSP, in many of these other cases, modifying the example so that the subject does intuitively seem to know the proposition in question also flips our intuition about the conditional that is crucial to the SCA account. As I noted, again and again SCA posits a certain block to our judging that we know, and the changes that would clear the way for our judging that we know also remove this block. This makes it difficult not to believe that SCA is at least on the right track. We will discuss these virtues of SCA in a bit more detail in section 5 of Chapter 6,

But perhaps there is another explanation to be had?

3. The Open Future: No Determinate Winner, Losers

One naturally imagines oneself into the lottery situation at a point in time when the winner has not yet been picked. (After all, after the drawing, we do seem to know we've lost.) So one might try to explain the difference in knowledge and in assertability between our two cases by appeal to the fact that there is not yet a determinate winner in the lottery situation. So it isn't determinately true that your ticket is a loser. So you can't know your ticket is a loser, since you can't know what isn't true. By contrast, there is a fact of the matter as to who won the Bulls game yesterday.

I have a good deal of sympathy for such thoughts, finding it very believable both that one cannot know things about the future that are not yet determinately true, and that in typical lottery situations, it is not determinately true before the drawing of any of the tickets that it will be a loser.

But while such rather general worries about our ability to know the open future might for some reinforce the appearance of ignorance in standard lottery situations, they cannot explain the more particular variety of apparent ignorance in play in lotteries, since that variety survives our moving the determination of the winner into the past. For even if the winner has already been picked in the lottery, so that there *is* now 1 winner and 19,999,999 losers, as long as the winning number hasn't yet been announced, the losers don't seem to know they're losers, and can't assert that they are. Some sweepstakes (at least profess to) work this way—'You may already have won.' Still, it seems, one doesn't know one is a loser.

To avoid complications involving whether one can know what isn't yet determinately true—complications that won't solve our puzzle anyway—let's stipulate that our lottery is one in which there already is a winning ticket (and many losers), but in which the winning number hasn't yet been announced. (Indeed, I've already here been putting the relevant judgments in the past tense, as concerning whether you know you have lost.) If you insist that there is no winning ticket until it has been announced (that it becomes a winner only at the announcement, not when the number is drawn), then alter the case so that the winner has been announced, but the people talking, though they know the announcement has been made, haven't yet heard what the winning number is.

4. The Existence of an Actual Winner: The Eccentric Billionaire's Lottery

Another type of explanation that might be initially attractive—in fact, a favorite of the person on the street—appeals to the claim that in the lottery situation, beyond the mere chance that your ticket is a loser, there is the actual existence of a winning ticket, which is in relevant ways just like yours. ('Somebody's gonna win.') By contrast, in the newspaper case, while there admittedly is a chance that your paper is wrong, we don't suppose there is an actual paper, relevantly like yours, which has the score wrong. This contrast is difficult to make precise, since, as I reported above, actual newspapers have indeed transposed scores. The claim must be that those newspapers aren't, in the relevant ways, like mine. Much depends upon which ways of resembling my paper are relevant. But on a fairly natural way of understanding that, only other copies of the edition I'm looking at are in the relevant ways like my copy of the newspaper, while just the other tickets to the lottery I'm playing will be like my lottery ticket. If so, then I won't think that there are other papers like mine in those relevant ways which have the score wrong, while I will think that there is a lottery ticket like mine in the relevant ways which is a winner. (We might then understand how our skeptic at the end of section 1 might start to get some traction in her skeptical urgings in terms of her trying to get you to think of a broader class of newspapers, which include some that have misreported scores, as being in the relevant way 'just like yours.')

Such an explanation can take several different routes at this point, but, it seems, any explanation that starts off this way is headed for trouble. For with many lotteries, there is no winning ticket. Many of the big state lotteries, for example, *usually* have no winner. Still, it seems, you don't know you've lost. In case you think that is because the jackpot is carried over to the next month's drawing, and then the next, and so on, until finally someone wins, so we think of the whole process as one giant lottery which will eventually have a winner, note that our ignorance of losing seems to survive the absence of that feature.

Suppose a somewhat eccentric billionaire holds a one-time lottery, and you are one of the 1 million people who have been given a numbered ticket. A number has been drawn at random from among 20 million numbers. If the number drawn matches that on one of the 1 million tickets, the lucky holder of that ticket wins a fabulous fortune; otherwise, nobody receives any money. The chances that you've won are 1 in 20 million; the chances that somebody or other has won are 1 in 20. In all likelihood, then, there is no winner. You certainly don't believe there's an actual winner. Do you know you are a loser? Can you flat-out assert you are a loser? No, it still seems. Here, the mere chance of being a winner—with nothing remotely like an assurance that there actually is a winner—does seem to destroy knowledge of your being a loser.

5. The 'Grabber' Lottery and Lewis's Account

The above case rebuts explanations that appeal to the claim that someone has won the lottery (so why not me?). However, while there is nothing like an assurance that there is a winning *player* or a winning *ticket* in our eccentric billionaire's lottery, it is part of its set-up that there is a winning *number*. Some explanations in the neighborhood we're considering might try to seize on that fact to handle the case.

The solution David Lewis advances at (Lewis 1996: 557), based on his rules of Resemblance and of Actuality, is an explanation in the neighborhood. On Lewis's account, one counts as knowing that p when one's evidence eliminates all the alternative possibilities to p that we are not properly ignoring. His rules concern the proper ignoring of alternative possibilities, the Rule of Actuality stating that 'the possibility that actually obtains is never properly ignored' (1996: 554), and the Rule of Resemblance dictating that if two possibilities saliently resemble one another and one of those possibilities may not be properly ignored, then neither may the other (1996: 556). When Lewis applies these rules to the lottery, he focuses on there being a winning ticket:

> For every ticket, there is the possibility that it will win. These possibilities are saliently similar to one another: so either every one of them may be properly ignored, or else none may be. But one of them may not be properly ignored: the one that actually obtains. (Lewis 1996: 557)

As such, his account may seem to be in trouble when applied to our eccentric billionaire's lottery, where there likely is no winning ticket, and yet we still get the characteristic lottery appearance of ignorance.

However, Lewis's account can perhaps be made to handle our case by focusing on the fact that there is an assurance in the set-up of the case that there will be a winning *number*. It can then be claimed that the possibility that my number will win saliently resembles actuality (that some other number wins), and so may not be properly ignored. Of course, much depends here on some fine points about salient resemblance. But we do well not to try to work that all out with respect to the example we are currently considering, because a variant of that case eliminates the feature of there even being a winning number, and yet still produces the characteristic lottery appearance of ignorance, and so should yield a more secure verdict.

So suppose again that our eccentric billionaire holds a one-time lottery, and you are again one of the 1 million people with a numbered ticket. This time, however, the mechanics of the drawing work differently. Not 20 million, but just one million balls, each with a number on it matching the number of one of the tickets have been placed in a giant vat, and thoroughly mixed around randomly. A button has been pushed which results in a mechanical 'grabber' being lowered into the vat, closed, and then raised up out of the vat. The grabber sometimes grabs a ball and raises it out the vat, but usually does not. Since the balls have been mixed around randomly, not only is there no telling whether any ball has been grabbed, but also which particular ball if any has been grabbed. Given the size and other properties of the balls and the grabber, there is a one-in-twenty chance that a ball has been successfully grabbed and raised out of the vat. (This has been exhaustively verified by many trial runs.) If a ball is grabbed and raised out of the vat, the holder of the ticket whose number matches the number on the raised ball wins a fabulous fortune; otherwise, nobody wins anything. So, once again, the chances that you've won are 1 in 20 million and the chances that somebody or other has won are 1 in 20. Do you know you are a loser? Can you flat-out assert you are a loser? No, it still seems. Here, the mere chance of being a winner—with nothing remotely like an assurance that there actually is a winning player, ticket, *or number*, nor that the possibility of your winning will resemble actuality in any salient way—seems to destroy knowledge of your being a loser. Suppose it is now revealed that the likely outcome did in fact obtain: the grabber failed to grab any ball. Now you know you didn't win, but we still judge that you didn't know this before that result was revealed to you, even though the possibility that you would win doesn't resemble actuality in any particularly salient way.

6. The 'Grabber' Lottery and Hawthorne's Account

Since we are discussing the 'grabber' lottery case, it's worth pausing to note the trouble it causes for John Hawthorne's account of the lottery intuitions. Hawthorne writes:

> Without pretending to be able to have a full account of the relevant psychological forces driving the relevant intuitions, we can nevertheless see that in the paradigm lottery situation, something like the following often goes on: The ascriber divides possibility space into a set of subcases, each of which, from the point of view of the subject is overwhelmingly likely to not obtain, but which are such that the subject's grounds for thinking that any one of the subcases does not obtain is not appreciably different than his grounds for thinking that any other subcase does not obtain. … In general, what is often at the root of the relevant lottery intuition is a division of epistemic space into a set of subcases with respect to which one's epistemic position seems roughly similar. Once such a division is effected, a parity of reasoning argument can kick in against the suggestion that one knows that a particular subcase does not obtain, namely: If one can know that that subcase does not obtain, one can

> know of each subcase that it does not obtain. But it is absurd to suppose that one can know of each subcase that it does not obtain. (Hawthorne 2004: 14-15)

But about our 'grabber' lottery, we judge we don't know that we've lost even though the features of our thinking about lottery cases that Hawthorne's account appeals to are absent: In thinking about this case, we do not divide epistemic space 'into a set of subcases with respect to which one's epistemic position seems roughly similar.' Rather, our grounds for the outcome we think is likely—that the grabber has failed to grab any ball and that therefore nobody has won—are very different from, and also much stronger than, our grounds for any of the other subcases into which we're inclined to divide epistemic space.

We should note that Hawthorne does not claim to be giving a 'full account,' and that he phrases his account in terms of how we '*often*' think concerning 'the paradigm lottery situation.' Nevertheless, one important way to test whether Hawthorne is really getting at the root of our usual tendency to judge that we don't know even in the paradigm lottery situation is to see what happens if we alter the case so that Hawthorne's account no longer applies. The 'grabber' lottery is such a test. That the characteristic lottery appearance of ignorance persists where the case is modified so that Hawthorne's account does not apply to it is reason to think Hawthorne's account is not correctly diagnosing where the appearance of ignorance is coming from, even in the paradigmatic lottery cases where the features he appeals to are present.

In a comparative vein, the 'grabber' lottery provides a nice test for deciding between SCA and Hawthorne's account. Of course, both accounts would lead us to expect that subjects will not seem to know that they (or others) will lose lotteries as those are standardly described. But, unlike Hawthorne's account, SCA predicts that we will still seem not to know we've lost in the grabber lottery, for about that case it will still seem that we would have believed we lost even if we hadn't lost. That gives us reason to think SCA, rather than Hawthorne's account, is getting at what's really driving the appearance of ignorance in lottery situations.[4]

---

[4] Note that Hawthorne has his own lottery case that he presses against SCA. Hawthorne points out that one typically *will* seem to know that one hasn't won a lottery in the different situation in which one isn't even playing the lottery—one doesn't even have a ticket. However, even the non-player's belief that she hasn't won seems to fail SCA's insensitivity test, for, as Hawthorne writes: '[I]f I had won, I would have owned a ticket, but not having heard the result yet, would think it a loser' (2004: 11). This case seems to favor Hawthorne's

7. The Existence of an Actual Winner: The Newspaper Lottery

In sections 4-6, we considered lottery cases that did not have as a feature of their set-up that there would be a winner. To approach this issue from the other side, what happens if there actually is a 'loser' newspaper? Suppose your newspaper announces that it has instituted a new procedure for checking and printing sports scores. This procedure has as a side-effect that one copy in each edition will transpose all the scores, reporting all winners as losers and all losers as winners, and, as there is no easy way for the distributors to tell which is the copy with the transposed scores, this copy will be distributed with the rest of them. But, as well over 1 million copies of each edition are printed, and as this new procedure will greatly cut down on the usual sources of error, this procedure will, on the whole, increase the likelihood that any given score you read is accurate. Here we've set up a virtual lottery of newspapers—one out of the one million copies of each edition will definitely be wrong. So we should expect our apparent situation vis-a-vis knowledge and assertability to match that of the regular lottery situation.

But put yourself in the relevant situation. You've heard about the new procedure, and so are aware of it. ('Good', you say. 'That means fewer mistakes.') Does this awareness affect your asserting practices with respect to the results of sporting events? I don't think so. You've read the newspaper, which is your only source of information on the game, and someone asks, 'Did the Bulls win last night?' How do/may you respond? I still say 'Yes, they did', as I'm sure almost all speakers would. I'd be shocked to learn that speakers' patterns of assertion would be affected by its becoming general knowledge that such practices, which, after all, increase reliability, are in place. As in the regular newspaper case, 'Probably' and 'It's quite likely that' seem quite unnecessary here in the newspaper lottery case. It still seems

account, for it happily seems not to apply to this non-player case. However, due to other proposed counter-examples to its insensitivity test, proponents of SCA, myself included, have refined that account of when it will seem that we don't know things. Hawthorne realizes this: Instead of just presenting his case and then declaring SCA vanquished, he sees that 'refinement is called for' (2004: 11) and pursues a refinement (involving methods of belief formation) due to Nozick before giving up on SCA. The problem is that he pursues the wrong refinement. The version of SCA I had already moved to in (DeRose 1995), due to very different challenging cases, already handles Hawthorne's example, as does a closely related version of SCA that Timothy Williamson bases on my account and pursues in (Williamson 2000a). But this will be covered in Chapter 6, sections 6-7.

you know they've won. Indeed, suppose that in this new case you are asked whether you know if the Bulls won. I respond positively, as I'm sure almost anybody would.

Of course, again, this appearance of knowledge may fade in the presence of a skeptic determined to make heavy weather of the possibility that your paper is the mistaken one. But your apparent knowledge that you have hands can also appear to fade under skeptical pressure. To repeat the point made in section 1 of this chapter, our current concern isn't whether, under pressure, one could be forced to retreat to 'Well, probably': that could happen in the *original* newspaper case. But as we *ordinarily* judge things, you *do* know the Bulls won in this newspaper lottery case, as is evidenced by your positive response to the question, 'Do you know?' and by your willingness to flat-out assert that fact when not under skeptical pressure. By contrast, we *ordinarily* judge, with no skeptics in sight (unless so judging makes *us* skeptics, in which case our puzzle is to explain why we're skeptics in the regular lottery case but not in the newspaper case), that we *don't* know we've lost the regular lottery, and that we can't assert that we have.

The newspaper lottery case combines elements of our two earlier cases—the regular newspaper case and the regular lottery case. Interestingly, with regard to one's belief that the Bulls won, the results in this new case match those of the regular newspaper case: you do seem to know, and can assert. Knowledge and assertability survive the actual existence of a 'loser' newspaper just like yours in the relevant respects. This, combined with the ability of our ignorance in the regular lottery case to survive the absence of a winning ticket, should put to rest the suggested explanation we've been considering that it's the existence of an actual winner that explains our difference.

8. SCA and the Newspaper Lottery

But the newspaper lottery's significance goes beyond the trouble it causes for that ill-fated explanation, which is one of SCA's rivals. The case provides this puzzle of its own. If one is thinking *only* about the newspaper lottery case, it seems pretty clear that we would continue to flat-out assert the results we've read in the paper, and would continue to think we know who won last night's games on the basis of having read them in the paper. But if one *compares* the newspaper lottery with the regular lottery, it can seem hard to reconcile that dictate about the newspaper lottery with the evident truth that we don't assert, and don't take ourselves to know, that we've lost a regular lottery. Isn't the newspaper lottery case *just like* the regular lottery? How, then, *could* there be this marked difference in our reactions?

Well, the newspaper lottery *is* just like the regular lottery in many relevant respects. But we should exercise caution in how we line the two cases up with one another in order to draw conclusions, or even expectations, from this similarity. What should this similarity lead us to expect? This, I submit: that, just as we judge that we don't know we've lost the regular lottery, so we will also judge in the newspaper lottery case that we don't know that we don't have the 'loser' newspaper. And this expectation is met: we do naturally judge ourselves ignorant of that fact. And that is just what SCA predicts, since we also tend to judge that one would believe that one didn't have the 'loser' newspaper even if this belief were false (even if one did have the loser newspaper).

In the newspaper lottery case, one will likely *believe* both that (a) the Bulls won; and that (b) I don't have the 'loser' newspaper. But, it is only the belief in (b) that SCA predicts we'll be blocked from counting as knowledge. One's belief in (a) escapes the block SCA posits, for we *won't* typically judge there that we'd now believe that the Bulls won even if they hadn't.

In the *regular* lottery, we judge that we don't know we've lost; this seems analogous to belief (b) in the newspaper lottery. What, in the regular lottery, is analogous to belief (a)? Well, suppose that I owe a friend a lot of money—so much that I am confident that I won't be able to pay off the loan by the end of the year.[5] Of course, I will easily be able to pay her back by the end of the year if I've won the lottery this week. Here, if I haven't yet heard what the winning numbers are, I'll likely believe both that (a') I won't be able to pay off the loan by the end of the year; and that (b') I've lost the lottery. While SCA correctly predicts that we'll think I don't know that (b'), my belief in (a') escapes SCA's wrath, since we won't typically judge that, in this situation, I would believe that I won't be able to pay up even if it were

[5] Against Sherrilyn Roush (2005: 132), note that we need not presume that *I won't be able to pay off the loan by the end of the year* entails *I won't win the lottery*. All I am presuming is that the comparative conditional *If I don't know that I won't win the lottery, then I don't know that I won't be able to pay off the loan by the end of the year* seems right.

the case that I'll be able to pay up. *Do* I seem to know, and can I assert, that I won't be able to pay off the loan this year? If asked whether I'll be able to pay up by the end of the year, while it is perhaps permissible for me to respond, 'No, unless I've won the lottery', it also seems perfectly permissible for me to answer with a simple 'No', not bothering my questioner with the remote possibility of my having won the lottery, just as I needn't bother her with the slight possibility that some multi-millionaire whom I don't know at all will pick my name out of the phone book this year as her sole legal heir just before dying.

So things look pretty good for SCA. It predicts that we won't think our beliefs in (b) and (b') constitute knowledge—and we don't. And our beliefs in (a) and (a'), which escape the block SCA posits, are beliefs we ordinarily would take to be knowledge. Of course, again, a skeptic can forcefully urge that we don't know, and shouldn't assert, that (a) or (a'), and he might even use our ignorance of (b) and (b') as part of her skeptical urgings. And, indeed, it is difficult to maintain that one knows that (a) (or (a')), while, in the same breath, admitting that, for all one knows, (b) (or (b')) is false. So we might well wonder whether we're right in naturally judging that we do know the former but not the latter. But these are all matters relevant to the issue of whether we really know—which issue we will get to soon. Our current concern is explaining the particular judgments that we would ordinarily make as to what we know and don't know, and what particular claims we'd typically be willing to flat-out assert if asked. And here, SCA gets things right.

9. What About 'My Paper is Accurate'?

You *believe* that your newspaper is accurate in the newspaper lottery case. But in that case, do you seem to *know,* and can you assert, that your paper is accurate when it comes to the sports scores it reports? Here we flip-flop. In settings in which we're focused on the fact that there is a 'loser' copy, we judge that we don't know this. In other settings, in which we're still perfectly well aware of the fact that there is a 'loser' copy but in which we're not particularly focussed on that fact, we may judge that we do know. I've been looking at scores from the paper I subscribe to for a long time, and I often come to have independent access to the results of the games it reports on. If it were inaccurate, in all likelihood I'd have known that by now. This would all remain true if my paper switched to a procedure which yields the newspaper lottery case. (In fact, that switch would make it more accurate.) Part of the reason for the flip-flop here may be an ambiguity in 'your newspaper'. Does this refer to the particular copy you hold in your hands, or to, say, *The Houston Chronicle,* a newspaper you and many others read every day?

Here SCA is supported by the fact that we similarly flip-flop on the subjunctive conditional SCA points us to. Where we're focused on the fact that there's a 'loser' copy, we're inclined to judge that you would still believe your paper was accurate, even if it weren't. To use the standard possible worlds analysis for subjunctive conditionals, this is because, given our then present focus, we take the closest world in which the antecedent is true (in which your paper is not accurate) to be a world in which the newspaper you subscribe to is generally reliable, but you happen to have the 'loser' copy. In this world, though your copy isn't accurate, you believe it is. In the other settings, in which you do seem to know that your paper is accurate, we take the closest world in which the antecedent is true to be a world in which the paper you subscribe to, say, *The Houston Chronicle,* frequently messes up. In such a world, your paper is not accurate, and you don't believe that it is, as you've noticed many of the frequent mess-ups. (At least this is true if, like me, you are a big sports fan who often looks at the scores and would have noticed if they were frequently wrong. If you are not thus like me, you may not seem to know your paper is accurate when it comes to its sports scores.)

10. Probabilistic Thoughts and Statistical Reasons

Addressing assertability in the lottery case, V.H. Dudman writes:

> It is not just that the probability is never high enough to trigger assertion. An exacter appreciation is that even the smallest uncertainty is enough to cohibit it. Assertibility goes out of the window as soon as the underlying thought is reduced to relying on 'mere' probability. (Dudman 1992: 205)

Dudman doesn't identify the probabilistic underlying thought involved in the lottery case, but, presumably, it is something like this: Only one ticket out of the 20 million is a winner; so, probably my ticket is a loser. By contrast, in either of our newspaper cases (the regular newspaper and the newspaper lottery cases), one's underlying thought is likely to be the non-probabilistic: The newspaper says the Bulls won; so, the Bulls won. Now, that the newspaper says the Bulls won doesn't entail that the Bulls won any more than there being only one winner out of 20 million lottery tickets entails that my ticket is a loser. But, plausibly, we in fact do tend to think probabilistic underlying thoughts in the lottery case but not in the newspaper cases.

Stewart Cohen has employed such a line of thought on our puzzle, attempting to explain why we don't seem to know in the lottery case, while we do appear to have knowledge in other cases much like my newspaper case. Cohen's account is couched in terms of the relevant alternatives theory of knowledge, according to which (at least in Cohen's hands) S knows that P if and only if S has a true belief that P and there are no relevant alternatives to P.[6] According to Cohen, while, in the cases in which we do think we know, there are alternatives to what we think we know which we're in no position to rule out, these alternatives are not relevant. By contrast, in the lottery case, we think we don't know precisely because we do find the alternative that we've won relevant, despite its great unlikelihood. What is crucial to Cohen's account of why we don't think our belief that we've lost the lottery is knowledge, then, is an explanation of why we find the chances of error *relevant* here, but not in the other cases. Cohen's answer is based on the 'statistical nature of the reasons' one has for thinking one has lost in the lottery case:

> What makes it [the alternative that one's ticket wins] relevant? I propose that the explanation lies in the statistical nature of the reasons. Although, as fallibilists, we allow that S can know q, even though there is a chance of error (i.e., there are alternatives compatible with his reasons), when the chance of error is salient, we are reluctant to attribute knowledge. Statistical reasons of the sort that S possesses in the lottery case make the chance of error salient. The specification that S's reason is the n-1/n probability that the ticket loses, calls attention to the 1/n probability that the ticket wins. (Cohen 1988: 106)

So Cohen's account ultimately is based on the 'statistical nature of the reasons' one has in the lottery case, which looks quite like Dudman's claim that we can't assert in

---

[6] According to most versions of the relevant alternatives theory, S knows that P iff (roughly) S has a true belief that P and can rule out all the relevant alternatives to P. But Cohen defines relevance in such a way that there can be no relevant alternatives to P where S knows that P (Cohen 1988: 101). Thus, in cases where S does know, what most versions of the relevant alternatives view classify as relevant but ruled-out alternatives are, for Cohen, irrelevant alternatives.

the lottery case because there our 'underlying thought is reduced to relying on "mere" probability'. Could either of these be the explanation?

Well, first, these accounts raise issues about the direction of explanation: Do we seem not to know because our grounds are probabilistic, or do we resort to grounds we phrase in merely probabilistic terms because we seem not to have knowledge-producing grounds for the items in question? And even if it is the case that we do think probabilistically and statistically in the lottery, but not in the newspaper case, and even if it were this statistical/probabilistic thinking/reason that blocked knowledge and assertion in the lottery case, we should want to know *why* we think merely probabilistic thoughts only in the lottery case*.* Perhaps all we think in the newspaper case (in typical cases, where, for instance, we're not confronted by a skeptic) is the non-statistical and non-probabilistic: The newspaper says it; so, it is so. But, we should want to know, why wouldn't analogous, non-probabilistic reasoning come off in the lottery case: It's a Super Lotto ticket (for heaven's sake!); so, it's a loser?

And, on the other side, it seems that in the newspaper (and the newspaper lottery) case, assertability, *pace* Dudman, can survive probabilistic thought. Consider the original newspaper case. Those few incidents in which my paper has transposed scores are often in the back of my mind when I rely on my newspaper for results of games. Still, I assert away. Suppose those incidents have worked their way to the *front* of my mind, as they sometimes do, as I'm asked, 'Did the Bulls win last night?' Suppose my underlying thoughts are consequently reduced to relying on 'mere' probability: 'The paper says the Bulls won; the probability that the paper's right is *extremely* high; so they probably won; it's overwhelmingly likely.' That is what I think. But what do/can I say? 'Yes, they won', seems just fine. 'Probably' seems quite unnecessary, despite the statistical nature of my reasons. My probabilistic thoughts and statistical reasons don't seem to rob me of assertability, or of knowledge. And why should they? Everyone should suspect that papers make occasional errors. Should this rob us of knowledge only when we're careful enough to think about it? And if I know the Bulls won, why can't I say they did? Indeed, I do say it, as would almost any other speaker. Assertability does *not* go out the window whenever the underlying thought is probabilistic. Assertability and knowledge can survive an abundance of merely probabilistic thought.

My conversational partner, after all, need neither know nor care whether probabilistic thoughts and statistical reasons happen to be guiding my thought at the moment. She wants to know if the Bulls won. Why should I trouble her with a 'probably' just because probabilistic thoughts and statistical reasons happen to be running through my head right now? If, on other occasions, where such thoughts are absent, I can flat-out assert, why should the fact that I happen to be privately entertaining such thoughts now affect how I should communicate with her? If she is well served by a simple, 'Yes, they won', on the other occasions, she'd be just as well served by that response now, and that, it seems, is what I should say. It *certainly* seems that I'm *allowed* to say it. Probabilistic thoughts don't block knowledge and assertability where we also have a way of knowing the item in question. So we would seem to need an account of why, in addition to being guided by probabilistic thoughts, we don't seem to know we've lost in the standard lottery situation—which is just what we've been seeking throughout the whole first part of this chapter.

So, in the newspaper case, assertability and knowledge seem capable of surviving probabilistic and statistical thoughts and reasons, while in the lottery case, we seem *somehow* prohibited from relying on the simple, non-probabilistic: 'It's a Super Lotto ticket (for heaven's sake!); so (of course!), it's a loser', to make the unqualified assertion 'I've lost' or to secure knowledge of our loss. So the suggestions we've considered in this and in the previous section can't explain the divergence in assertability and in apparent knowledge between our cases. Again, SCA can.

11. Causal Connections

The attempted explanations we considered in the previous section echo an element of Harman's own proposed solution in (Harman 1968), where Harman denies that the statistical grounds available to one in a standard lottery case can give one knowledge that a ticket has lost (1968: 166). However, for Harman, this claim is underwritten by a deeper account of what is going on in our (well, really, his) cases. Even if we agree that knowledge 'goes out the window' when our 'underlying thought is reduced to relying on "mere" probability', that only raises the question of what's so reducing us to such a reliance. Harman has an answer. He claims that all inductive inferences take the form of inference to the best explanation (1968: 165), and these seem to be inferences to the best *causal* explanations. The picture that results seems to be this. In the newspaper case, your belief that the Bulls won, or the evidence you can access to base such a belief upon, is in part caused by the fact that the Bulls won in a way that allows you to perform a good inference to the best causal explanation to the conclusion that the Bulls won, and thereby come to know that fact, even though the causal connection here is not reliable enough to yield a perfectly airtight inference. (Presumably, though, it has to be at least fairly reliable to be knowledge-producing.) In the lottery case, you've heard about how the lottery works, and we suppose that you are thereby causally hooked up in a reliable enough way with the various facts about the lottery's set-up that make it very likely that you have lost, so, by a good (though not perfectly airtight) inference to the best causal explanation, you can know those facts. And, for Harman, you can also know what you can deduce from those facts. But such a deduction does have to be airtight: When you reach beyond what you can access by an inference to the best explanation, you can do so knowledgeably only by airtight deduction. And all you can so *deduce* from those facts about the lottery's set-up is that you have *probably* lost, not that you have lost, so you can't know the latter. Harman's account in (1968), then, is based on a causal restriction on knowledge: You cannot know you've lost because your losing has no effects that you can access as evidence by which to infer that you've lost via an inference to the best explanation. And indeed, my understanding is that Harman was working on that paper while Alvin Goldman was working on his famous 'A Causal Theory of Knowing' (Goldman 1967), and each knew about what the other was up to.

As we know from the discussion of causal theories of knowledge, sometimes we know empirical facts, not because they cause our beliefs in them or the evidence on which such beliefs can be based, but because both those facts and our beliefs in or evidence for them have a common cause. Indeed, here we don't have to reach beyond what was already contained in Goldman's classic paper; he noted this at (1967: 364). Accordingly, as I argued at (DeRose 1996: 569-70, n. 3), it seems that when my copy of the newspaper reports the Bulls as winning, I can know, not only that the Bulls won, but also, apparently pace Harman, that my neighbor's copy of the same newspaper reports them as winning. My belief that my neighbor's copy reports that result is not caused by the fact that it does, and the fact that it does cannot be reached by the kind of inference to the best causal explanation that Harman seems to demand if one is to know a fact by way of a non-airtight inductive inference. Rather my neighbor's copy's having the result that the Bulls won and my evidence for thinking that it does have a common cause in the fact that the Bulls did win (and they presumably have other common causes a bit 'downstream' from the Bulls' win as well).

The obvious fix is for the causal theorist to loosen up his account so that all that is required is that one's belief and the fact that it is a belief in be 'causally connected' in a way that allows for knowledge in cases where they have a common cause—and indeed, that's just what Goldman does. But then, as David Christensen points out, the causal solution to the lottery puzzle is undermined on the other side.[7] For if we allow knowledge of facts by being so 'causally connected' in a reliable enough (even if not perfectly reliable) way with those facts, where we allow that such a connection can be one where one's belief or evidence and the fact it's a belief in have a common cause, then it seems that in lotteries of the type we're considering, with enough tickets and few enough winners, I should be able to know that I've lost. Lotteries like that, after all, are *extremely* reliable (even if not perfectly reliable) producers of losing tickets, and we're already supposing that our beliefs concerning the basic set-up of the lottery are produced in a way that allows us to know those facts. So, it looks like the lottery system produces losing tickets in an extremely reliable way, and also can produce our beliefs to the effect that we're losers. On the loosened-up causal account, then, it becomes unclear why we shouldn't think that I know that I have lost the lottery.[8]

[7] This doesn't undermine causal theories of knowledge. Some such causal theorists might be happy to accept that lottery losers do know that they have lost. They would then need a damage-controlling explanation for why that can seem otherwise, but they would not be alone in that: As we saw in section 1, everybody has some damage control to do. But this does undermine the causal account of the two judgments about the individual cases. When a causal theorist of knowledge gets into the game of giving an account of those two judgments, they then enter the market for an account of those judgments. Their causal theory of knowledge won't by itself provide the account; what kind of damage control they propose will be a crucial part of it.

[8] Citing my example, Christensen similarly uses a case where a belief and the fact that is its object have a common cause against a proposal in (Nelkin 2000) that is in important ways like Harman's, and Christensen then notes the possibility of the causal theorist moving to an account that only requires a causal connection (in a way that allows for common causes), and then argues, in the way I have followed in the text above, that this move would undermine the causal solution to the lottery puzzle (Christensen 2004: 61-2).

12. That There is a Chance of Winning is the Whole Point of the Lottery!

Lotteries may be extremely reliable producers of losing tickets, but it's not as if they *aim* to produce only losers. It's important to the whole enterprise that there be the odd winner. That's a feature; not a bug, as they say. Might that be important to the appearance of ignorance we are trying to diagnose? Here we reach a certain kind of proposed explanation that some readers may have been inclined toward ever since our puzzle was presented in section 1. I won't spell out this explanation in full; I suspect that it can be completed in various significantly different ways. But the various explanations I have in mind are all based on the observation that, with respect to the belief that one has lost the lottery, the chance that this belief is wrong—i.e., the chance that one is a winner—is intimately connected to the whole point of entering the lottery. By contrast, in the newspaper case, while there is a chance that your belief that the Bulls won is mistaken, this chance is not similarly connected to any of your goals. It's just an unwelcome and unintended side-effect of or bug in the process by which you come to have your information.

The newspaper lottery in section 7 may have reinforced this suggestion in some readers' minds. In that new case, though we've set up what in many ways is a lottery-like situation, we retain knowledge and assertability. Why? Because, the suggestion under consideration goes, having a 'loser' newspaper is not any part of the point of the new procedure. It is still just an undesired side-effect. Knowledge and assertability in this new case match that of the old, regular newspaper case, and diverge from the regular lottery case, because, like the regular newspaper case and unlike the regular lottery case, the chance of your being wrong in our new case is not correctly connected to any relevant goals.

But the reflections of section 8 should show us why such a suggestion cannot provide the explanation we've been seeking. It is only when we focus on your belief that the Bulls won that you seem to know and can assert in the newspaper lottery case. But if we instead focus on your belief that you don't have the 'loser' newspaper, you seem to lack knowledge and assertability. Here, the chance that you are wrong (i.e., the chance that your copy is the 'loser') does seem to prevent you from knowing, despite both its minuteness and its buggy lack of a connection with any relevant goals.

13. The Big Pay-Off, etc.

Closely related to the proposed explanation explored in the last section is this slightly different, but equally doomed, proposal. It can be tempting to think that it is the great pay-off one will receive if one has won the lottery that justifies us in treating seriously, despite its minute probability, the possibility that one has won—or, even if it doesn't *justify* our so treating that unlikely possibility, it at least explains why we *do* so treat it. But this can't be our explanation, for our apparent ignorance in lottery situations survives the absence of a big pay-off, as the reader can quickly verify by considering how assertability and apparent knowledge would fare in a lottery with no pay-off at all—one held 'just for the fun of it'.

The following lottery-like example will further illustrate the ineffectiveness of this explanation, together with a host of other explanations built upon various observations regarding our goals and interests which I won't take the space to investigate one by one. Suppose you learn that one copy of a phone book with a great circulation—say, the *Greater Houston White Pages*—contains, in its printing of the second 'f' of George T. Jefferson, III's name, ink of a different type from the ink used in the rest of the phone book. Although you've learned this fact, you are *completely* uninterested in it. Nobody else finds it interesting either. Even if there were an easy way of discovering whether your copy is the one with the differently inked 'f', you wouldn't lift a finger to find this out. Despite your complete lack of interest in the matter, it will still seem to you that you don't know that yours isn't the copy with the 'strange' 'f'. You'll seem every bit as ignorant here as you are of your not being the winner in the lottery case, where your interest in whether you are the winner, and the pay-off involved if you are the winner, is great. So your ignorance in the lottery case seems not to stem from anything having to do with big pay-offs, our interests, and the like. It seems the accounts under consideration can't solve our puzzle. Again, SCA can.

14. Our SSP Solution Applied to the Harman Lottery Puzzle

Up to this point, we have been concerned with explaining the two particular judgments constitutive of the Harman lottery puzzle. I have defended SCA as the correct account of this. Now we turn to the comparative judgment I've claimed should be thought of as another (a third) piece to our puzzle—and to the issue of how best to actually solve this puzzle.

The most straightforward way to incorporate SCA into an account of knowledge is of course to hold that our concept of knowledge just *is* that of sensitive true belief. On this view, we tend to think that insensitive beliefs, like that we've lost the lottery, aren't known because they simply aren't known, and we recognize that fact, while we also correctly recognize that we do know the likes of *The Bulls won*, about which our beliefs are sensitive. We could then join Harman (and many others) in endorsing both of the intuitive judgments about the individual cases. But such an account runs into the same problem that Harman himself faces when we apply it to the comparative judgment concerning the lottery cases: Like other accounts that simply deliver the intuitively correct verdicts about the individual cases, its verdict concerning the comparative matter is as counter-intuitive as the individual judgments it verifies are intuitive. Any such account owes us some good damage control concerning the comparative judgment here, and, absent that, it is highly dubious.

But fortunately, as we saw in SSP, there is a not-so-straightforward way to use SCA in a contextualist account which can actually *solve* the puzzle of skeptical hypotheses—and can solve the lottery puzzle, too. Indeed, if you're willing to count *I've won the lottery* as a skeptical hypothesis, the latter solution can be counted as an instance of the former. We can add another pair (though I'll add two not-Hs and two Os in this case, the idea being that we can pair either of the not-Hs with either of the Os) to the bottom of our chart of 'epistemologically perplexing pairs of propositions' from section 10 of SSP:

not-H	O
I'm not a BIV	I have hands
Those animals aren't just cleverly painted mules	Those animals are zebras
The paper isn't mistaken about whether the Bulls won last night	The Bulls won last night.

I've lost the lottery/ I haven't won the lottery[9]	The Bulls won last night/ I won't be able to pay off my loan this year

As with the top three pairs, we can again sense that the following comparative fact holds for our new lottery pairs: I am in no better a position to know that O than I am in to know that not-H. This comparative fact is revealed in each case by the highly plausible conditional, *If I don't know that not-H, then I don't know that O*. Yet, as was also the case with our old pairs, when we consider our new Os by themselves, they seem like things we know to be the case, but when we consider the not-Hs, there is at least a strong tendency to say/think that we don't know them to be the case. And our beliefs in these Os seems sensitive, while any belief we might have in these not-Hs seems insensitive. All of this indicates that we can apply our solution from SSP to our lottery puzzle.[10] (The main difference between it and our solution to the puzzle

[9] We should consider the lottery puzzle as formulated in terms of 'I haven't won', both because that is the kind of formulation it is usually given in the literature, but also to allow us to consider Hawthorne's variation on the example, in which one does not even have a ticket, which we will see in section 6, and answer in section 7, of the following chapter. (Where one doesn't even have a ticket, it seems that you do know you haven't won, as Hawthorne points out. This is an important example to account for. However, you won't seem to have *lost*, since you just aren't playing, and so Hawthorne's point would be blocked on the 'I've lost' formulation.) But we should keep the 'I've lost' formulation as well, to consider in connection with possible solutions to our puzzle like the one pursued in (Nagel 2011). Key to Nagel's psychological account of why we seem not to know that we haven't won the lottery is that we tend to engage in the more careful 'System 2' thinking when we consider that matter (while we would tend to utilize the more easy-going 'System 1' style of thought when considering the likes of our Os), and one of the main 'triggers' that induce the shift-up to System 2 thought is the kind of sentential negation involved in the 'I haven't won' formulation of our puzzle, but not in the 'I've lost' formulation. There are other triggers that can be in play (including, crucially, presentation of numerical odds (Nagel 2011:11)), but since the one involving negation seems very important, the 'I've lost' formulation of the puzzle can provide some key tests for solutions of Nagel's type. See esp. (Nagel 2011: 10-18), for related discussion.

[10] In the first paragraph of sect. 12 of SSP, I present my proposed 'Rule of Sensitivity' in the way it appeared in my dissertation (DeRose 1990): making use of the notion of comparative strength of epistemic position, but without appealing to possible worlds. (In my dissertation, I just used 'comparative conditionals' to get a fix on comparative strength of epistemic position.) Then in the second paragraph of that section, I proceed to explain the notion of strength of epistemic position and the Rule of Sensitivity in possible worlds terms. Here I apply my solution to the AI skeptical puzzle to the lottery puzzle, but in its original form, without the explanation in terms of possible worlds. The original form transfers nicely, but the later elaboration raises some tough questions when applied to the lottery puzzle. It involves supposing that the possibility that I have won the lottery is very distant. But while we do speak of such possibilities as 'remote' (and the like), lending hope to the thought that something in the vicinity of our possible worlds story might apply here, the type of

the other pairs confront us with is that a particular problem, that we will see in section 19 and deal with for most of what follows that section, emerges for our solution when we apply it to the lottery puzzle.)

The key to the solution is the account of why insensitive beliefs seem to us not to be instances of knowledge, developed in sections 11-12 of SSP. The SSP-like solution that I now propose to the Harman lottery puzzle adopts that account.

The comparative judgment so troublesome to other treatments of the lottery puzzle is no problem at all for the SSP-like solution I am proposing: That's the moving part that we nail down first and most securely (in the way that we nailed down a similar comparative judgment in section 10 of SSP). We simply accept the comparative judgment that we are in at least as good an epistemic position with respect to *The Bulls won* in the newspaper case as we are with respect to *My ticket is a loser* in the standard lottery case. We hold that the conditional, *If I don't know that my ticket is a loser in the lottery case, then I don't know that the Bulls won in the newspaper case* (and also its contrapositive), evaluated on comparative grounds, is true, no matter how high or low the epistemic standards are set (so long as they are not allowed to vary from one to the other instance of 'know' in the conditional), in the same way that in SSP (again, sect. 10) we held that *If I don't know that I'm not a BIV, then I don't know that I have hands* is true at any epistemic standard. And similarly, about the choice between the 'weird' and the 'normal' lottery ticket we considered in section 1, we are not saddled with saying that in any good sense of 'know', you should there choose the ticket you know will lose over one that you don't know will lose.

But if we are in as strong a position with respect to *I've lost the lottery* as we are with respect to *The Bulls won*, why, when we consider the scenarios individually, are we inclined to judge that we don't know the former, but that we do know the latter? It's here we adopt the explanation employed in SSP (sects. 11-12) for why insensitive beliefs seem not to be instances of knowledge. On the resulting account, insensitivity is given the starring role in our explanation of why we seem not to

---

remoteness in play here seems to be a matter of the improbability of the hypothesis, which improbability seems to result from how thinly we are slicing up possibilities, and so seems not to be something well captured in terms of how distant is the nearest possible world in which the possibility is realized. A notion of safety in which it isn't measured simply by how close is the nearest danger, but also takes into account the density of the nearby dangers, might be the way to go here.

know that we've lost the lottery—the starring role that, in the previous sections of this chapter, we saw that it deserves. Importantly, however, insensitivity plays this role not because sensitivity is taken to be part of the very concept of knowledge, but instead because of the crucial role it plays in the operation of the mechanism—the 'Rule of Sensitivity'—by which the content of knowledge attributions is manipulated so that claims that someone 'knows' something that they don't sensitively believe tend to go false.

Such an account can thereby *solve* the Harman lottery puzzle: It can avoid *simply* endorsing the particular judgments that it nevertheless explains, while allowing us to strongly endorse the comparative intuition that proves so troubling to other views.

Applying the account to the Harman lottery puzzle yields a view on which *I've lost the lottery* gets the same treatment that in SSP we gave to *I'm not a BIV*—a treatment on which in standard lottery situations losers *do* know that they've lost, *by ordinary, moderate standards for knowledge*. Lottery winners, of course, don't know that they've lost, even by such moderate standards, but they would be rational to take themselves to know that they have lost by ordinary standards for knowledge. Our tendency to think we don't know we've lost the lottery is then accounted for by the Rule of Sensitivity, on which, because our belief that we lost is an insensitive one, saying that one does or doesn't 'know' such a thing will tend to raise the epistemic standards to high levels that the belief does not meet, so that admissions that we don't 'know' will tend to come out true, while claims to 'know' such a thing will tend to be false.

On this account, we tend to think we don't know that *I've lost the lottery* while we think we do know that *The Bulls won*, not because we simply don't know the former while we simply do know the latter,[11] but because a claim to 'know' the former, by

[11] Here we use 'simply' in the semi-technical way explained at (DeRose 2009: 228-9): S 'simply knows' that P iff S is such that any speaker who is using 'knows' standardly, no matter their context, would be asserting a truth if they were to say that S 'knows' that P. This could be either because invariantism is true and S meets the non-varying requirements for knowledge with respect to P, or because contextualism is true, but S meets all the standards for knowledge allowed by the standard use of 'knows'. Similarly, then, S 'simply doesn't know that P' iff S is such that any speaker who is using 'knows' standardly, no matter their context, would be asserting a truth if they were to say that S 'doesn't know' that P.

the Rule of Sensitivity, invites a much more demanding reading than does a claim to 'know' the latter—an understanding of 'know' so demanding that we don't satisfy it.

15. The Standard Contextualist Solution to the Harman Lottery Puzzle

As I did in section 2 of SSP with respect to the skeptical puzzle, we can also speak of a 'basic contextualist strategy' to the Harman lottery puzzle that my account is an instance of, but which can also be taken by other contextualists who don't appeal to the notion of sensitivity as I do. We can call the resulting type of treatment the 'standard contextualist solution' to the Harman lottery puzzle, where this designates any contextualist account on which, in standard lottery situations, we do know that we've lost the lottery by ordinary standards for knowledge, but which explains why we seem not to know by claiming that talk of whether we know we've lost the lottery tends *somehow* (whether through the operation of a Rule of Sensitivity, or by some other mechanism) to raise the epistemic standards to high levels that our belief that we've lost does not meet. Cohen's (1988: 106-111) and Lewis's (1996: 565-6) treatments of the lottery seem also to utilize this basic strategy.

I do not myself think the standard contextualist solution is tenable (or even provides a real solution to the puzzle) when it is not taken at least roughly in the insensitivity-based direction I take it, but since there are prominent instances of the basic strategy that, at least by my lights, do go in such wrong directions, it is worth noting when a problem or consideration affects the basic strategy in general, and not just my particular implementation of it.

16. The Intuitive Pull (Felt by Some) Toward Judging that We Do Know that We've Lost the Lottery

Toward the end of section 1, I wrote that the particular intuitions we'd be busy trying to explain in the first half of this chapter are 'rather flighty and dodgy in various ways.'[12] I think almost everybody can feel the intuitive pull of these verdicts—though no doubt to varying degrees. And I hope to have provided a good account of when and how it is that that intuitive pull operates on us. However, many (including in my experience some who are generally very thoughtful and discerning about such matters) feel pulled in opposite directions when considering lottery cases, also feeling, in some other part of themselves, as it were, an opposing pull toward saying that maybe we losers really *do* know we've lost the lottery—just as many are somewhat torn about the first premise of our skeptical argument (AI), feeling some significant pull toward insisting that they really do know that they're not BIVs. And I'm sure that at various points in our discussion, especially discerning readers had worries about just how intuitive are the judgments that we don't know that we've lost in various lottery situations that I was laboring so hard to explain.

I must tread lightly here, since individual reactions vary so much. I want to be cautious about claiming an advantage for our contextualist solution in being well positioned to explain why there might also be a significant intuitive pull toward thinking that we do know we've lost the lottery, for some might not feel any such pull that needs to be accounted for (so, for some, providing for such an explanation will count *against* a view), while even those who do feel such a pull can differ with one another over its nature and so over what would well account for it.

However, I think some who do feel a significant intuitive pull toward judging that we do know we've lost the lottery will find our contextualist solution quite congenial in being able to account for that pull—just as many who feel a significant pull toward thinking they do know that they're not BIVs may feel a special attraction to SSP's solution of the skeptical puzzle because it is able to account for that pull (see section 9 of Chapter 2). For, after all, on our solution, a loser's belief that they've lost the lottery has a high degree of what it takes to count as a piece of 'knowledge'. You are as well-positioned in the way relevant to knowledge with respect to *I've lost the lottery* as you are with respect to *The Bulls won*. The reason a

---

[12] The intuition about the comparative issue can also be something one feels conflicted about. Of course, whenever one faces a puzzle consisting of jointly inconsistent intuitions, there will be potential for one to feel conflicted about each of the intuitions constitutive of the puzzle, since its falsehood will be implied by other things one finds intuitive. Here, that one is inclined, when the cases are considered individually, to judge that one does know that the Bulls won, but that one does not know that one has lost the lottery, can of course push one to doubt that one really is in as good a position with respect to the latter as to the former.

claim to 'know' the former tends to go false while a claim to 'know' the latter tends to come out true is not that your belief in the latter has more of what it takes to be 'knowledge', but rather only that we tend to set the threshold for what will count as 'knowledge' higher when we address the former. Though our belief that we've lost the lottery doesn't meet the epistemic standards that, due to the operation of the Rule of Sensitivity, we tend to apply *to it*, on our solution it is a belief that we are very well positioned with respect to, and it does meet the ordinary standards for knowledge that we most typically employ *in general*. In light of those facts, it would seem the reverse of surprising that some of us would feel some tendency toward judging that we do know that we've lost a standard lottery.

17. Ordinary-Strength Claims to Know that Someone Has Lost the Lottery: The Case of Andy, Nico, and Lou

On a related note, our solution can handle special conversational situations where it seems that one *can* appropriately and truthfully say that someone in a standard lottery situation (in which their relevant belief is insensitive) 'knows' that a player (perhaps themselves) has lost the lottery.

Recall that on our contextualist solution, claims to know propositions that one is well-positioned with respect to but that one only insensitively believes tend to be false, not because they fail some standing sensitivity requirement for knowledge, but due to the operation of a conversational standards-raising mechanism, the Rule of Sensitivity. Consequently, we should expect there to be situations where other conversational forces are in play, pushing the epistemic standards in directions different from where the Rule of Sensitivity is pushing them, and perhaps prevailing over the Rule of Sensitivity. This could result in claims to the effect that subjects 'know' things that they only insensitively believe that nevertheless seem or even are appropriate and/or true. (We shouldn't expect these appearances to be extremely clear or forceful, though: Even if some other conversational force overcomes that exerted by our Rule, this will be a case where the 'conversational score' is being pushed in conflicting directions, and the Rule of Sensitivity will still be at least muddying the waters.) I have long thought that there were situations in which one could with propriety and truth claim that someone 'knows' that a skeptical hypothesis is false, though one's belief that the hypothesis is false is insensitive. My favorite such situations are ones in which speakers are describing subjects as 'knowing' such things in order to compare the ordinarily strong epistemic position those subjects are in with respect to the propositions in question with the deplorable epistemic position some other unfortunate subjects may be in with respect to those same things (sometimes because these unfortunates don't even believe the propositions). We will look a bit more closely at these situations in section 8 of the following chapter. But if, as I'm claiming, *She lost the lottery* works much like *I am not a BIV*, we should expect that analogous situations should produce similar claims that someone knows that they lost the lottery. So it is worth quickly noting here that that expectation is met.

To convert one of the examples,[13] suppose Andy devises a practical joke that a group of fun-lovers attempt to play on their friends, Nico and Lou, on which they will try to get those two to believe that the ticket they jointly own to the big state lottery has actually won, and that they are the co-winners of a fabulous fortune. After Andy temporarily leaves the scene, Nico catches someone snickering, and has to be let in on the joke, but Lou is completely taken in by it. So, when Andy returns, one of the friends pulls him aside to apprise him of the situation, explaining: 'Nico knows that they haven't really won the lottery. Sorry, we had to tell her. But Lou totally bought it!' Here Nico, like all her friends, does not sensitively believe that she has not won the lottery. Yet here it seems that 'knowledge' that she hasn't won is quite naturally and appropriately—and I think truthfully—attributed to her. The relevant intuitions about these cases may be a bit delicate, but whatever verdicts one wants to issue about the status of this claim that Nico knows that she hasn't won the lottery, the situation should seem to you very similar as that of the analogous claim that Nico knows that the animals in the zebra cage aren't cleverly painted mules that we'll consider in the analogous case in section 8 of Chapter 6.

[13] Based on the relevant example in section 8 of Chapter 6, the reader can also easily construct for herself an analogous case involving a psychiatric ward of patients all of whom believe they have actually won a big lottery. You should end up with a claim that the psychiatrist's assistant 'knows' that he didn't win the lottery which to a significant extent seems quite appropriate and perhaps even true.

18. Ordinary-Strength Claims to Know that Someone Has Lost the Lottery: 'Come Off It' / 'Get Serious' Claims

Another variety of claims to the effect that someone knows that someone has lost the lottery that can seem appropriate are what we can call 'Come off it!' claims. Williamson addresses some such claims here (the bracketed portion is in brackets in the original):

> There is a special jocular tone in which it is quite acceptable to say '[Come off it—] Your ticket didn't win', but the tone signals that the speaker intends not to make a flat-out assertion. (Williamson 2000a: 246)

We should note that it seems similarly acceptable to throw 'You know' in and say, with a similar tone, '[Come off it—], You know your ticket didn't win.'

However, though my own example in the previous section of an acceptable claim to know one has lost the lottery does involve a practical joke, I don't think 'jocular' well captures the tone with which many such remarks would be said—including the remark in my case. Williamson has strong theoretical motivation for so-characterizing the claims in question, since he's committed both to the knowledge account of assertion (as I am) and also (as I am not) to the view that one *simply* doesn't know that one didn't win the lottery. If the sayings in question were acceptable as assertions, that would be trouble for Williamson. Calling the tone involved 'jocular' promotes the thought that these claims are not serious flat-out assertions for which his theory is accountable.

And indeed, insofar as such claims require a special tone for their acceptability, and don't seem acceptable when said 'flat-footedly', I think we have good reason to be cautious in our use of them.[14] And *perhaps* saying such things in a 'jocular' tone is one way to make them acceptable—though imagining them said in a joking fashion isn't really helping *me* to hear them as acceptable. (Though my example in the previous section involves a practical joke, what is said about Nico wouldn't seem to be something that would be said in a jocular tone: That seems to be just a piece of straightforward information conveyed to Andy to seriously apprise him of the jocular situation.)

---

[14] For relevant discussion, see (DeRose 2009: 15-18), (DeRose 2009: 97-98, n. 20), and (DeRose 1998: 70-2).

In fact, the way for the 'Come off it!' claims to be said that helps me to hear them as acceptable often is the *reverse* of imagining a jocular tone: They often seem most at home when said in a very serious, '*Stop* joking around!' way. The tone involved is earnest, insistent, and *dismissive* of frivolity. Indeed, I think the same category could also be well labeled 'Get Serious!' claims. For example: Alice is in financial peril, but is trying to avoid and/or delay asking her parents for help in paying the rent due at the end of the month, and suggests she should hold off because she might have won the big lottery whose winner will be announced and paid the day before her rent is due. 'Come off it! You know you didn't win the lottery,' you might well say. And whatever else may be going on with such a scolding, that's no joke.[15]

Though I proceed here with due caution, the contextualist approach to the lottery puzzle would seem to provide a nice way of dealing with such lottery claims. While Williamson holds that, in the sense we've been using the phrase, one *simply doesn't know* that one has lost the lottery (in the standard lottery situation), the contextualist solution posits that, by ordinary, moderate epistemic standards, one does know that one has lost the lottery, and the appearance of ignorance here is explained by the claim that attempts to claim such 'knowledge' tend to invoke the high epistemic standards that one does not meet, and so *tend* to backfire, while admissions that one doesn't know that one has lost the lottery tend to come out true. The very serious claims that subjects 'know' that individual players have lost that we are here discussing can then be understood as claims where the speakers do not, or intend not to, raise the standards in the way that claims with such content normally would. A function of the insistent tone with which such claims tend to be made is perhaps in part to signal such an intention.

---

[15] Compare Hawthorne's discussion at (Hawthorne 2004: 18-19). Citing Williamson's remark, Hawthorne grants that 'we sometimes make knowledge claims using a tone indicating that we are not to be taken literally,' but then immediately goes on to write: 'But I see no good evidence that this is always going on in these cases' (2004: 18).

19. Hawthorne's Objection and Multi-Premise Closure

We turn now to the key objection against solutions like ours, leveled by Hawthorne.[16]

Consider the contextualist's claim that you know$_m$ that you've lost in a standard lottery situation, where 'knows$_m$' designates knowing by the moderate standards by which, according to the contextualist, you do know such a thing.[17] But note that this contextualist account of your belief that you have lost applies as well to beliefs you might have about other individual players of the same lottery: You will also know$_m$ that Alfred lost, and that Betty lost, etc. Indeed, we can number the individual ticket holders (or the individual tickets, if you prefer), and this contextualist view is that you can know$_m$ of each of them that they have lost—except, of course for the winner: about that player, you can reasonably take yourself to know$_m$ that she lost, but you don't actually know$_m$ that, since your belief that she lost isn't true.

Trouble ensues for this contextualist account when you start drawing inferences from all of these supposed bits of knowledge$_m$. (And this is trouble for the basic contextualist strategy, and not just my particular way of implementing it.) Supposing that you know$_m$ the conclusions that you deduce from premises you know$_m$ to be true, you should on the contextualist view know$_m$ that *none* of you, Alfred, or Betty has won. But it's one thing to say that you know$_m$ of *each* of you losers that she is a loser; it's quite another to say that you know$_m$ that you're *all* losers. In a big enough lottery, that may be alright when we stick to just the three of you, but of course we don't have to stop there. As we saw, the contextualist holds that you can know$_m$ of each of the losers of the lottery that they lost, so we can build bigger and bigger groups of losers, and the contextualist seems committed to your being able to know$_m$ even of these big groups that *all* (and not just each) of the

---

[16] Hawthorne's argument is at (Hawthorne 2004: 94-5), where it is aimed at standard contextualist solutions generally; he gave the argument specifically against Lewis in his earlier (Hawthorne 2002).

[17] Of course, on contextualist analysis, there are no doubt many different standards for knowledge that are 'moderate' and that often govern ordinary talk. This should not affect to points to follow, where we employ the handy fiction that we have identified a particular set of moderate standards that we are designating by 'knows$_m$.'

players in them have lost. At some point, that commitment becomes bizarre, even when keeping firmly in mind that we are talking about knowledge by only moderate standards here. When the groups get big enough—say, by the point that they include over half of the players—it seems pretty clear that you just don't know by *any* allowable standard, and so don't know$_m$, that none of the players in the group has won. And things get *really* crazy when we consider the group of *all* the losers, since even here the contextualist seems committed to saying that you can know$_m$ that all of them have lost. And since we can suppose you also know$_m$ the facts about the lottery set-up needed to fuel this deduction, as Hawthorne nicely points out, it looks like you're then in a position to deduce, and thereby come to know$_m$ of the winner that's she's won! But that's truly absurd: To you, that winner looks just like all the other players—like someone who is almost certainly a loser.

Note that this argument works not only against solutions that use what I am calling the 'basic contextualist strategy', but against any position on which it is held that one knows by a particular standard, or that one simply knows, that one has lost in the usual lottery situation—at least when this is held for reasons that would also apply to one's beliefs that other individual players have lost. For a variety of reasons, some of them quite respectable, an invariantist might conclude that in the typical lottery situation, you (if you are indeed a loser) really do know that you have lost. But if this invariantist holds this for reasons that will also apply to your beliefs about other losers of the lottery, she too will be in danger of falling prey to Hawthorne's attack.

But note also that Hawthorne's argument relies on some form of *multi-premise closure*, the principle that in its general form lays it down that, as Hawthorne likes to formulate it:

> Necessarily, if S knows $p_1, \ldots, p_n$, competently deduces q, and thereby comes to believe q, while retaining knowledge of $p_1, \ldots, p_n$ throughout, then S knows q. (Hawthorne 2004: 33)

We should not overstate what is required for an argument of Hawthorne's type. Though this would not suffice to drive his opponents to full height of absurdity reached two paragraphs above, Hawthorne could force an opponent to a conclusion that is, to my thinking, quite sufficiently bad, making use only of repeated applications of some more limited principle, like two-premise conjunction:

> Necessarily, if S knows both p_1 and p_2, competently deduces their conjunction, p_1 *and* p_2, and thereby comes to believe the conjunction, p_1 *and* p_2, while retaining knowledge of both of p_1 and p_2 throughout, then S knows the conjunction, p_1 *and* p_2.

But a Hawthorne-like argument does require *some* form of at least somewhat controversial multi-premise closure, and this seems to be the premise of his argument most vulnerable to attack. And this may in turn explain why his argument appears to be particularly effective against the standard contextualist solution to the lottery puzzle (and perhaps why Hawthorne aims his argument in that particular direction). For the contextualist solution seems to be motivated in the first place by some kind of commitment to closure. It is a distinctive move of standard contextualists—both in treatments of skepticism and lotteries—to explain away apparent failures of closure as instead being due to context shifts. The standard contextualist claims that closure holds, but not for *the* knowledge relation (since there is no single knowledge relation), but rather for any of the particular knowledge relation*s* that get invoked by uses of 'know(s)'—like knowledge$_m$, with which we are currently working. Hold that relation constant, according to standard contextualists, and closure holds up: You know$_m$ that you have hands, but you also know$_m$ that you're not a BIV; you know$_m$ that you won't be able to repay your loan by the end of the year, but you also know$_m$ that you've lost the lottery. The reason you can seem not to know that you're not a BIV or that you've lost the lottery is that talk of whether you 'know' such things tends to invoke unusually high epistemic standards that you do not meet. Though BIV talk would typically invoke standards much higher than that summoned by lottery talk, for our current purposes it does no harm to pretend that the same unusually high standards, those governing 'knowledge$_h$,' say, are put into play in the cases, allowing us to say that you know$_h$ neither that you're not a BIV nor that you've lost the lottery, but that closure is upheld here, too, for you similarly know$_h$ neither that you have hands nor that you won't be able to repay your loan. Given the allegiance to closure paid by standard contextualists, we seem particularly vulnerable to Hawthorne-like arguments, which utilize some form of multi-premise closure as their only controversial premises in driving their opponents to quite unacceptable conclusions.

20. Toward Intuitive Closure: Problems and Refinements

But even the best friends of closure realize that it is difficult to formulate an acceptable closure principle, because annoying problems keep popping up when you try to do so. Still, there can at least seem to be a fairly clean division between supporters and opponents of closure. Here is perhaps the most prominent opponent of closure, Robert Nozick, admitting the intuitive power of the closure principle, which he is calling 'P,' after noting some difficulties of formulation:

> We would be ill-advised, however, to quibble over the details of P. Although these details are difficult to get straight, it will continue to appear that something like P is correct. (Nozick 1981: 205)

This sets up Nozick's denial of closure in his next paragraph:

> Principle P is wrong, however, and not merely in detail. (206)

Nozick wants to be clear that he's not just nit-picking here, but denying the very spirit of closure—even without an adequate statement of the principle's letter. Affirmations of closure tend to have the same character. Consider, for instance, this prominent pro-closure statement—the emphasized sentence of which has become a fairly well-known anti-anti-closure salvo—by Richard Feldman:

> Roughly, the closure principle says that if you know one proposition, and know that that proposition entails another, then you know the latter proposition. There are details that might worry us about this, and some philosophers, notably Robert Nozick, have denied the closure principle. *To my mind, the idea that no version of the closure principle is true—that we can fail to know things that we knowingly deduce from other facts we know—is among the least plausible ideas to gain currency in epistemology in recent years*. But I won't argue that here. For the most part, I will just assume the truth of some version of the closure principle. (Feldman 1999: 95, emphasis added)

Feldman is affirming closure in just the way that Nozick, whom he explicitly cites, denies it. He may not have a correct formulation of what it is he's affirming, but he is claiming—quite emphatically—that *something* in the vicinity of his rough statement of closure is right. And similarly, though just in a footnote (#33), in SSP I say that I, unlike Nozick, do believe that 'something like P is correct.'

But it's easy to worry about these denials and affirmations of the spirit of closure, where we don't have a correct statement of its letter. Might there not be issues that come up where it isn't clear whether something is really part of the spirit of closure? If we call the holy grail here—the correctly refined formulation of closure that best captures the intuitive spirit of the principle—'intuitive closure,' might there not be disputes about what should be included in intuitive closure, even before we get to the issue of whether it's right or wrong? Might it be problematically unclear what the dispute over closure is about?

These questions are especially pressing for me here, because, although I'm often classified as a defender of the closure of knowledge (indeed, some seem to think of me as an *arch*-defender of it), some recent defenses of closure, notably by Hawthorne and Williamson, suppose that 'intuitive closure,' as we're using the phrase, includes aspects I think need to be weeded out. And, perhaps unsurprisingly, one of the refinements that I think is needed would derail Hawthorne's argument against the contextualist solution to the lottery puzzle.

So I will argue in what follows that we should accept a form of closure that is weakened in a way that undermines Hawthorne's argument, even while, as I aim to show in Appendix D, it remains of use to the skeptic who utilizes an AI-like argument. But to see the proposed limitation in the much-needed proper perspective, we should start by taking a look (a very brief one should do) at the project of trying to arrive at a properly formulated closure principle.[18]

To get something of the flavor for the project of refining ('chisholming'[19]) the closure principle, start by considering it in its roughest form. Even Feldman's rough statement already included an important refinement. An even rougher and simpler principle (and one even more desperately in need of refinement) would be a statement like this (whose wording mimics Feldman's):

> If you know one proposition, and that proposition entails another proposition, then you know the latter proposition.

---

[18] For a much more thorough look at the issues that come up in trying to correctly formulate the closure principle, and for ways of dealing with those problems, see (David & Warfield 2008).

[19] Definition from *The Philosophical Lexicon* (http://www.philosophicallexicon.com/): 'chisholm, v. To make repeated small alterations in a definition or example. "He started with definition (d.8) and kept chisholming away at it until he ended up with (d.8'''''''')."'

Imagine someone finding that principle quite plausible. But then, we can easily suppose, a nasty problem suddenly occurs to them: 'Wait a minute! The mere fact that the first proposition, which I know, entails some other proposition doesn't guarantee that I'll know that second proposition. For all that, I might not even be aware of the entailment!' So they might refine the principle to Feldman's 'rough' statement of it (presented here with added emphasis):

> If you know one proposition, *and know that* that proposition entails another, then you know the latter proposition.

But, as Feldman realizes (which is why he says this formulation is still rough), further problems come up. For instance: 'Wait a minute! The mere facts that I know something and know that it entails something else doesn't guarantee that I'll know that something else. For all that, I might not even believe[20] the second proposition! Maybe I know both the first proposition and the fact that it entails the second, but I've never put those two bits of knowledge together to infer the second proposition. Perhaps I even deny that second proposition.' So you try further refinements.

I bring up these preliminary problems with formulating closure mostly to point out one of their potentially important characteristics. Note that these are problems that occur in the abstract: We didn't have to say what the 'first proposition' and 'second proposition' (or the 'p' and the 'q', if that's how you formulate things) are to raise the problems. It's not that the form of the principle that one is currently considering seems right until you apply it to some particularly problematic propositions, where you see that it doesn't always seem to work.

---

[20] Keeping with the policy I used in the first volume (DeRose 2009: 186-7), I am pretending in this book that belief is the 'attitude of knowledge', and so here ignoring potential problems with closure that might result from belief not being that. Suspending that pretense for a moment here in this note, where certain closure principles require that the subject believe the conclusion of her deduction before they rule that she knows that conclusion, I will be interested in modified versions of those principles that instead require that the subject have the attitude of knowledge toward her conclusion. Connected with that, I'll note that the example which drives Krista Lawlor's case against closure principles strong enough to underwrite skepticism is one in which the subject she considers fails to be confident enough to know the conclusion of his deduction (Lawlor 2005: 32-4). If we construe our closure principles as requiring that the subject have the attitude of knowledge toward his conclusion where the principles we are considering explicitly require belief in the conclusion, Lawlor's case should be no counter-example to our principles so construed. I suspect that those who think that belief *is* the attitude of knowledge will think that Lawlor's subject does not really believe his conclusion, and so does not constitute a counter-example to the unmodified closure principles that include a belief requirement.

Rather, these are issues that can arise before we even start considering particular propositions that the closure principle might be applied to.

We now skip some stages in the process (which can go in some interestingly different ways), and get to Hawthorne's formulation, which bears the marks of the problems that happened to come up, at least as he apparently thought through the project. As will soon become very important to our concerns, note that in getting to where Hawthorne is, we didn't just keep limiting the principle further and further to take care of problems. Somewhere in there we also had the (excellent) idea that the intuitive closure we seek applies not just to instances where we deduce a conclusion from a single known premise, but also to cases where we deduce a conclusion from multiple premises that are all known, and so we *broadened* the principle (allowing deduction from a single premise as a special case of our broader principle).[21] After a few tweakings, we arrive at Hawthorne's formulation of multi-premise closure, which I repeat here:

> Necessarily, if S knows $p_1, \ldots, p_n$, competently deduces q, and thereby comes to believe q, while retaining knowledge of $p_1, \ldots, p_n$ throughout, then S knows q. (Hawthorne 2004: 33)

21. Yet Another Problem: The Aggregation of Risk

But now a still further problem might well occur to some, as it did to me: 'Wait a minute! What if, for each of those premises, there's a risk, from S's point of view, that it's wrong? We're supposing S knows all these premises, so that means supposing that they are all in fact true, and that the risk of error from S's point of view is in each case small enough that it doesn't preclude S from knowing the premise. But suppose all those risks *are* small enough to allow for knowledge. Maybe they're truly tiny, but since there may be a multitude of premises involved here, if all the premises are needed for the deduction (and let's suppose they are), all those little bits of risk could together lead to trouble. When S puts all those premises

---

[21] Well, in many ways it would have been better for our search to have begun from a rough version of multi-premise closure. I started our tour instead with a rough version of single-premise closure mainly to make contact with Feldman's formulation.

together and deduces the conclusion, can't the resulting risk (from S's point of view) of being wrong with respect to the conclusion be too great to allow for knowledge of it?'

Note that, though this new scruple may perhaps be more advanced than some of the others, the need for further modification here is still seen in the abstract: It is not a trouble that comes up only when we apply the principle to some particular propositions, but rather one that we can see as a problem without any such specification. Why should this not be added to the long list of problems calling for tweakings of the closure principle?

I raised such a problem, calling it the problem of the 'accumulation of doubt' at (DeRose 1999: 23, n. 14; DeRose 2004c). Hawthorne raised and discussed the problem, phrasing it in a better way (that we will follow here) as a problem of the 'aggregation of risk,' at (Hawthorne 2004: 46-50), before in the end, maintaining his allegiance to a form of multi-premise closure that does not correct for the problem. And the problem was more recently pressed by Maria Lasonen-Aarnio (2008), who ably defends it as a genuine problem.[22]

Now, one might think that this is a problem just for multi-premise closure, since in the case of an inference from a single premise, there are not multiple premises to serve as multiple sources of risks that can then aggregate. And indeed, Hawthorne frames his discussion as being of a potential problem that threatens only multi-, and not single-premise, closure (2004: 47). But I think this is a mistake. Even when there is only a single premise, one is still relying on both the premise and its connection to the conclusion when one makes a deduction. There can be some shakiness in S's deduction of the conclusion from that premise that can join forces with the shakiness that the premise has for S to perhaps occasion enough shakiness in S's epistemic grasp of the conclusion to prevent her from knowing it, even having deduced it, at least fairly competently (and at a level of competence that often produces knowledge of the deduced conclusion), from the premise.[23] Indeed, when I

[22] I like Hawthorne's wording better than my own, old phrasing, because, as I found myself adding when explaining the problem in talks and in classes, the problem needn't be one of actual doubts that the subject is feeling, but can just concern whether the subject *should* feel some doubt. Lasonen-Aarnio wisely follows Hawthorne's wording here.

[23] Cf. Lansonen-Aarnio: 'In short, the problem with SPC [Single Premise Closure] is the following. When a subject comes to believe a proposition *Q* solely based on competent deduction from *P*, her epistemic standing

discussed the problem at (1999: 23, n. 14), it was explicitly in the context of a discussion of single-premise closure, pressing just this concern. Given my sense of how risk, in the relevant form, aggregates (on which it works at least much like the notion of the *probability* of being wrong, from a subject's vantage point), and some other assumptions, I suspected, and still suspect, that failures of single-premise closure would be cases in which one just barely knows the premise, and just barely knows that it entails the conclusion (or, using Hawthorne's formulation, where one's deduction is just barely competent),[24] and one then falls just barely short of knowing the conclusion. (And that's how I presented my worry about single-premise closure in the work cited.)

---

with respect to *Q* will depend *both* on her epistemic standing with respect to *P and* on the competence of her deduction. But because competence doesn't require infallibility, the risk involved in her belief in *P* and the *deductive risk* involved in the deduction itself can add up so that despite satisfying the antecedent of SPC, the subject fails to satisfy its consequent, and fails to come to know *Q*. Call these *accumulation of risk* failures of single premise closure' (2008: 159-160).

[24] At (DeRose 1999: 23, n. 14), I was working with a rough formulation of closure which (like the formulation Feldman presents) contains a clause specifying that S knows that the entailment holds. There is some reason to instead formulate the principle in the way Hawthorne has it. As Lasonen-Aarnio points out (2008: 161-2), we might well want our principle to be applicable to subjects who don't have the concept of entailment, and so don't know that entailments hold, but who can be competent deducers nonetheless. Now, one could specify a meaning for the semi-technical term of 'competent deduction' that would prevent this problem from arising for single-premise closure, by not counting any episode of deduction that is *at all* shaky for the subject as a 'competent' one. (But in that case, 'competent' seems far too mild a word for what one has in mind.) But one could do the same thing with formulations that instead specify that S knows that the entailment holds by beefing that up to the condition that S knows *with absolute certainty* (with no shakiness whatsoever) that the entailment holds. And such a principle could be interesting for various purposes. (For instance, as we will see in Appendix D, skeptical arguments might work well with such a formulation.) And indeed, you could get an interesting principle of multi-premise closure that heads off this problem and that starts off: 'If S knows that p_1, and also knows *with absolute certainty* each of $p_2, \ldots, p_n$, and super-competently deduces …' But if we are looking for refinements to the closure principle that don't *over*-solve for the problems that come up, things get trickier. I'm assuming that we're requiring knowledge-level security of deductions (that the security the deduction has in one's hands is at the level provided by knowledge of holding of the entailment: one doesn't have to know that the entailment holds or have the concept of entailment, but the deduction must be as secure for one as it is for those who do know that the entailment holds) for them to count as 'competent' when I express my suspicion that failures of closure here that are due to aggregation of risk will have to be 'just barely' failures. If 'competent' is used in another way, all bets are off; if it's used without somehow specifying the level of security required, then I don't sufficiently understand the resulting principle.

22. Fixing the Closure Principle to Address the Problem of the Aggregation of Risk Undermines Hawthorne-Like Objections

My doubts about single-premise closure may be controversial (though see also the discussion of such doubts in (Lasonen-Aarnio 2008)), but that multi-premise closure can fail for reasons of aggregated risks seems at least fairly clear, doesn't it? And such failures needn't involve any 'just barely's: *Multi*-premise closure can fail *badly*.[25]

That seems fairly clear *to me*. But even clearer (to the point that it perhaps deserves to be uncontroversial) is the meta-judgment that it is far indeed from clear that multi-premise closure is perfectly fine without containing any fix for this problem. Thus, I think anything deserving of the name 'intuitive closure' will contain a 'fix' for the problem of the aggregation of risk. We will look at what this weakening of the closure principle might look like in Appendix D, where we will see that the weakening does not jeopardize the skeptic's use of closure in AI. But we can see from the outset, before getting into all the details, that once we accept the need for such a weakening of the closure principle, the resulting modified principle will no longer underwrite Hawthorne-like objections against our solution to the Harman lottery puzzle, for these objections rely precisely on ignoring the very problem in question as they try to saddle our solution with unwelcome commitments to subjects having remarkable bits of knowledge$_m$; the deductions these objections propose are just the sort that any such fix will be designed to warn us off of.

To review, our solution claims that losers in standard lottery situations do know by ordinary, moderate standards (that is, they 'know$_m$') that they have lost the lottery—and others can know$_m$ this about them, and losers can know$_m$ this about each other. (Crucially, we also give an explanation of why things can seem otherwise, and an account of what to say about the winners of lotteries, but for now we focus on the aspects of the solution that the objections in question seize upon.) Hawthorne-like objections seek to discredit our solution by claiming that, given the

[25] *Multiple applications of* single-premise closure can also fail badly, i.e., can take you from a starting point that is known even by quite demanding standards, to an ultimate conclusion that falls far short of meeting even quite lax standards for knowledge.

knowledge$_m$ that our solution posits, we will be committed to our being able to know$_m$ by deduction various things that we clearly do not know$_m$. ‘Well, if you know$_m$ that you have lost, then you should also know$_m$ that Alice has lost, and know$_m$ that Betty has lost, and …,’ these objections begin—correctly (supposing none of the players named is the actual winner). We say about all these bits of knowledge$_m$ just what we say about your knowledge$_m$ that you yourself are a loser: You do (or at least can) know$_m$ of each of these losers that they lost, and we can explain why things might seem otherwise. But the objection continues: ‘And from all this knowledge$_m$, you should be able to deduce, and thereby come to know$_m$, that …’ And here it is alleged that our solution commits us to the presence of knowledge$_m$ of the types we’ve considered (of problematically large groups of losers that they are *all* losers, and of the winner that she’s won) where clearly no ‘knowledge’ of any kind exists. But we can now see that these objections here rest with all their weight on an overly-strong form of the closure principle that does not include a proper fix for the problem of the aggregation of risk. So we reply to the objections that though we do know$_m$ that the various losers are losers, we cannot come to know$_m$ the various problematic things that the objection claims we’re committed to the knowledge$_m$ of, because the deductions in question all rely on many premises, that, though they are known$_m$, are such that there is some risk that we are wrong with respect to them, and when we draw conclusions from that multitude of known$_m$ premises, the risk from our point of view that we are wrong with respect to the conclusion is (often far, far) too great to allow for knowledge$_m$ of those problematic conclusions.

23. An Infallibilist Evasion of the Problem?

It seems clear that if there is some risk of being wrong that attaches to some of the known premises of a competent deduction, the deducer can in the relevant cases fail to know the conclusion when and because, due to the aggregation of risk, there is too great a risk, from her point of view, of that conclusion being false. How might advocates of a strong form of multi-premise closure that doesn’t contain a fix for this problem of the aggregation of risk (and that can therefore be used to drive Hawthorne’s argument; we can call such strong forms of closure ‘Oxford closure’) respond? Defenders of Oxford closure might do well to cast a critical eye upon the thought that there might be some risk from the subject's point of view that the various premises are false, if in fact she knows them to be true. The idea would be to save Oxford closure from the problem of the aggregation of risk by claiming that when the premises of a deduction are really known, then there simply is no risk—none whatsoever—of being wrong there to be aggregated. This could be called an attempted 'infallibilist' evasion of the problem.

Though this evasion of the problem seems based on an implausibly strong infallibilist notion of knowledge, it can appear to have a reasonable basis in how, in our talk and thinking, we seem to connect the notion of knowledge with that of risk—and with other notions nearby to that of the risk of being wrong. Here I have in mind various forms of tension involved in claiming to know that p while, at the same time, admitting that there is a risk that you might be wrong about p—or that there is a *chance* or a *possibility* that one is wrong, or that one *may* or *might* be mistaken. Conjoining a claim to know with such an admission, all in one breath, certainly seems to result in an unhappy utterance; 'I know that p, but there is some risk that p is false', for instance, would seem to qualify as a 'clash' of the type discussed in volume 1 (DeRose 2009: esp. 96-8, incl. notes and 208-9, n. 17). And when one's admission that there is some risk (or chance, or possibility) that one might be wrong is still hanging in the conversational air, it can seem at least awkward, or perhaps wrong, to claim to know the proposition in question—and vice versa. In various ways, the rules of the knowledge/risk language game can seem to require that one *choose* between saying that one knows something and admitting that there is a risk (or chance or possibility) that one is wrong about the matter, as it can seem wrong to stand by both. And one might take all this to show that our notion of knowledge is tied to those like the risk of error in such a way that the knowledge of something completely precludes there being any risk whatsoever of one's being wrong about the thing in question.

24. Micro-Risks of Error and the Failure of the Infallibilist Evasion

But surely that *is* an implausibly strong infallibilist construal of knowledge! What's more (though *perhaps* this is just unpacking what's meant, or often meant, or what one often intends to mean, by such phrases as 'no risk whatsoever'), this evasion of the problem requires, not just that there be nothing that is properly called a 'risk' of error in cases where 'knowledge' is correctly attributed, but that there not even be any 'micro-risks' of error, as we might call them.

To explain: Consider Unger's old comparison of 'knows' with 'flat'.[26] In the previous section, we briefly considered some tensions between claiming to 'know' something and admitting that you suffer from some 'risk' of error about the matter. (I explore some ways of handling such tensions in Appendix F.) Unger noted quite similar tensions between calling a surface 'flat' and admitting that it had 'bumps'.[27] Largely from such tensions, he concluded if a surface is flat, then it has no bumps whatsoever. However, Unger claimed, just about all the physical surfaces we encounter, including those we typically describe as 'flat', do in fact have some bumps, however small, and so, he argued, they are not actually flat:

> For example, while we *say of* many surfaces of physical things that they are flat, a rather reasonable interpretation of what we presumably observe makes it quite doubtful that these surfaces actually *are* flat. When we look at a rather smooth block of stone through a powerful microscope, the observed surface appears to be rife with irregularities. And, this irregular appearance seems best explained, not by its being taken as an illusory optical phenomenon but, by our taking it to be a finer, more revealing look of a surface which is, in fact, rife with smallish bumps and crevices. (Unger 1975: 65)

Unger concluded that we are just about always speaking falsely when we call an ordinary physical surface 'flat'.

Those hoping to cling to the truth of our ordinary claims that surfaces are 'flat' have several options (and we won't come close to touching on all of them here) as to what to say about Unger-like arguments and the microscopic 'smallish bumps' he points out. Employing our semi-technical use of 'simply', I suppose that one could bravely declare that the 'irregularities' revealed by microscopes *simply* are *not* bumps, and anyone who calls them 'bumps' in any context is speaking falsely. Those things simply aren't big enough to really constitute 'bumps'! Alternatively, one could

---

[26] (Unger 1971); updated in Chapter 2 of (Unger 1975)

[27] Unger usually worked with the adjective 'bumpy', rather than the noun 'bump', so as he would usually put things, he was exploiting tensions between calling a surface 'flat' and admitting that it was 'bumpy', though he did on occasion (as in the quotation I'm about to display in the text) slip into talk about the noun.

agree with Unger that it turns out that almost all the physical surfaces we talk about, including those we typically call 'flat', simply do have bumps, but hold that being flat, and being truthfully called 'flat', is consistent with having (small enough) bumps. But many will be tempted by a contextualist approach to 'bumps', saying that what counts as a 'bump' varies with context. In most contexts, the microscopically small 'irregularities' Unger points to don't rise to the level of counting as 'bumps'—which is why we do (despite our standing knowledge of what the surfaces of physical objects are like), and can, truthfully say of many ordinary physical surfaces in many contexts that they have 'no bumps'. (Or even, in some perfectly natural sense, 'No bumps whatsoever': You worked for a long time sanding a tabletop to remove all the large bumps, but still left about ten 'small bumps,' as you were happy to call them. The next day, you sand those down. 'Now there are no bumps whatsoever,' you say, apparently [or at least this seems apparent to me] truthfully, despite the continued presence of Unger-bumps.) Relative to those contexts, the small 'irregularities' that Unger writes of are what we might well call 'micro-bumps': they are things that don't count as 'bumps' relative to the context under discussion, but do count as 'bumps' in other contexts, where they meet the lower thresholds in place for what is to count as a 'bump'. Such a contextualist line would seem to hold promise for accounting for the power (such as it is, and this does seem to vary much from person to person) of Unger-like arguments for flatness skepticism. Perhaps pointing out such small irregularities and calling them 'bumps' is a way of putting into place, or at least of pushing toward putting into place, standards by which they do count as 'bumps'—of making what were just micro-bumps now count as 'bumps'. And here one can sense possibilities (that we won't here explore) for explaining the tensions that Unger exploits in ways that won't involve us in any jarringly incredible claims to the effect that almost all our positive ordinary claims to the effect that ordinary physical surfaces are 'flat' are false.

It's easiest to make sense of a notion of 'micro-bumps' when one accepts contextualism about 'bump', so one can say, as I do above, that a micro-bump relative to context c is something that does not count as a 'bump' in c, but does count as a 'bump' in other contexts with more liberal standards. But suppose that for some reason (perhaps an irrational aversion to context-sensitivity) one is an invariantist about 'bump', and so holds that, though the standards for what we are likely *to call* a 'bump' may well vary a lot among contexts, what can be *truthfully* called a 'bump' does *not* vary from context to context. Well, then, you won't think there are any micro-bumps relative to any contexts—at least given our *current* understanding of 'micro-bump'. But it seems you might still have good use for a related notion, which might well be given the same label. For, supposing you're not a *skeptical* invariantist like Unger,[28] but instead think that the standards for what counts as a 'bump' hold steady at some *moderate* level (at which microscopic 'irregularities' certainly don't count), you will often face things much like those you would count as 'bumps' but which you'll think don't—and that sometimes don't *quite*—rise to the level of counting as a 'bump', and you might have use for a term like 'micro-bump' to describe such things, where by this you will mean roughly: something much like a bump, but which doesn't rise to the level of being a bump; something that, if there were only more of it, would be a bump.[29] If we can make sense of such a notion (and I admit, it's not the clearest in the world), that should become our general notion, available to both contextualists and invariantists about 'bump'. For our new characterization, we can use the explication just given, but relativize it to contexts: A micro-bump, relative to context c, is something which is much like the things that count as 'bumps' in c, but which doesn't rise to the level of counting as a 'bump' in c; it is something that, if there were only more of it, would count as a 'bump' in c. For invariantists, the line between 'bumps' and micro-bumps does not move from context to context. For contextualists, it does. This gives the contextualist an added aid in answering questions of the likes of 'Wha'dya mean, "something that, if there were only more of it, would count as a bump in c"?!': for the contextualist can often add that it is something that *does* count as a 'bump' in contexts other than c, which are more liberal than c in counting things as 'bumps'.

Though it's not a physical notion of something we can visually imagine, we can utilize an analogous notion of 'micro-risks' of error, understanding them to be micro-bumps in the epistemic road. A 'micro-risk' of error, relative to context c, is

---

[28] Our 'skeptical' invariantist here is actually maximally *liberal* in what she'll count as a 'bump'; we call her 'skeptical' because, given certain views of the relation between 'bump' and 'flat', she is likely to be maximally stingy in counting things as 'flat'.

[29] So micro-bumps will often 'behave' in ways similar to bumps. So, to use an Unger-inspired illustration, just as you should be mindful of the potential effects of bumps in the field you're playing on when planning a shot in croquet, or when you're predicting how likely your planned shot is to succeed, so, especially when it's an important shot and one that calls for much precision, you might do well to consider what some prominent micro-bumps might do to your shot. (We're supposing here that 'you' are a moderate invariantist. Contextualists, for whom the boundary between bumps and micro-bumps is flexible, will likely hold that where speakers have interests that make what were just micro-bumps very important to the conversation, the boundary will shift so as to make those things count as full-fledged bumps.)

something which is much like the things that count as ‘risks’ of error in c, but which doesn’t rise to the level of counting as a ‘risk’ of error in c. It is something that, if there were only more of it, would count as a ‘risk’ of error in c. This notion and/or closely related notions (micro-possibilities and micro-chances of error) may be helpful in characterizing what ‘infallibilism’ about knowledge is: Perhaps we can say that the infallibilist, but not the fallibilist, holds that knowing that p is incompatible with there being risks *or even micro-risks* (and/or micro-possibilities, and/or micro-chances) of error with respect to p? And perhaps our fallibilist will be able to explain the tensions we looked at in the previous section in ways that won’t involve us in any incredible claims to the effect that even the tiniest micro-risks of error are enough to make even ordinary claims to ‘know’ something go false? I’ll explore these matters a bit in Appendices E and F. But for our current purposes, what is important to note is that micro-risks of error certainly seem to be just the kind of thing that, when they join forces in great enough number, can aggregate into full-blown, knowledge-destroying risks of error. To recall (from note 29) Unger’s illustration: As even micro-bumps in a field can really mess up a croquet shot when there is a whole multitude of them, if you’re relying on enough premises in drawing a deduction, and there are big enough micro-risks of error attaching to enough of the known premises you are relying on, then it seems you can fail to know (and even fall *far* short of knowing) the conclusion you have competently deduced because there is from your point of view too great a risk that you are wrong about that conclusion for you to know it (all the while keeping the standards for knowledge constant).

Thus, to make good on the evasion we considered in the previous section, it looks like one will have to insist that there cannot be even the tiniest micro-risks of error with respect to the premises of a deduction if those premises are really known. (And, indeed, something like that is what one may well be trying to convey, perhaps successfully, when saying the likes of that, for the infallibilist, there cannot be ‘any risk at all’, or ‘any risk whatsoever’, or ‘absolutely any risk’ of error when something is really known.) And that certainly does look like an implausibly strong infallibilism about knowledge.

25. The Infallibilist Evasion and Standard Contextualist Solutions to the Lottery Puzzle

The infallibilist evasion of the problem of the aggregation of risk is especially problematic when one is ultimately targeting our standard contextualist solution to the lottery puzzle. (Recall from section 19 that Hawthorne-like objections can also target invariantist approaches to the puzzle.) For, recall, our claim is explicitly just that we losers *know*$_m$ that we have lost the lottery, and it would certainly seem that knowledge *by ordinary, moderate standards* does not require a total absence of even the tiniest micro-risks of error! Indeed, it can seem part of the very notion of moderate, as opposed to very stringent, standards for knowledge that the former allow for knowledge in the face of small enough risks or micro-risks of error. One is in fact tempted to say that even unusually stringent epistemic standards allow for very small micro-chances of error, and it is at most only on *maximally* stringent, absolute standards (perhaps usually in play only in certain philosophical discussions of questionable value) that knowledge can be incompatible with even the very tiniest micro-risks (micro-chances, micro-possibilities) of error.

26. Intuitive Closure and Oxford Closure

As I noted in section 19, standard contextualist approaches to both skepticism and the lottery have involved a strong allegiance to (at least the spirit of, if not any particular formulation of the letter of) closure, and as I noted in section 20, I in particular am often thought of as a defender, and even an arch-defender, of closure. So some readers will be surprised, not to mention a bit suspicious, to learn that I seek to escape a key objection to my solution to the lottery puzzle by denying a form of closure (what I'm calling 'Oxford closure') that has been prominently endorsed, and even, as we are about to see, called 'intuitive'. That is why it was important for me to put the issue in its proper perspective, and see the needed weakening of the closure principle as just one more of the many tweaks that we all know perfectly well such principles need to endure to be truly plausible.

Williamson gives the name of 'intuitive closure' to this principle:

> Deduction is a way of extending one's knowledge: that is, … knowing $p_1, \ldots, p_n$, competently deducing *q*, and thereby coming to believe *q* is in general a way of coming to know *q*. (Williamson 2000a: 117)

And he defends the principle mainly by appeal to its intuitive attractiveness:

> We should in any case be very reluctant to reject intuitive closure, for it *is* intuitive. If we reject it, in what circumstances can we gain knowledge by deduction? (2000a: 118)

And Williamson seems to be using 'in general' as one might use 'in full generality'—in a way that precludes exceptions (or that at the least precludes exceptions for cases of aggregation of risk). This makes it easier to understand his challenge to specify the circumstances under which we can extend our knowledge by deduction. His pointed challenge: If you don't think this principle holds in full generality, then *you* tell *me* when it is that we *can* extend our knowledge by deduction! The answer, in its salient part, is: when the problem of the aggregation of risk doesn't prevent the extension; see Appendix D, at n. 12 and the text that note attaches to, for what closure principles that contain a fix for this problem might look like. Beyond that, the formulation of the closure principle is no more our problem than his. On this reading of 'in general', Williamson's 'intuitive closure' is strong enough to drive Hawthorne's argument. But on this reading, 'intuitive closure' is no longer intuitive. Viewed from the proper angle, it seems quite intuitive indeed that one could fail to extend one's knowledge over deduction due to an aggregation of risk. At the very least, a principle that rested with all its weight on the impossibility of such failures would have at most an extremely shaky claim to being 'intuitive'.