Australasian Journal of Philosophy 74 (1996): 568-580
This should be close to how the paper appeared in AJP, but will not include the
final changes that were made to it. Please quote only from the published version.

KNOWLEDGE, ASSERTION, AND LOTTERIES

Keith DeRose

I. The Problem(s)

In some lottery situations, the probability that your ticket's 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're 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 quite alright to say, but, it seems, you're 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's a loser, and (b) In flat-out asserting some proposition, you represent yourself as knowing it.[1] This working hypothesis will enable me to address these two phenomena together, moving back and forth freely between them. I leave it to those who reject the hypothesis to sort out those considerations which properly apply to the issue of knowledge from those germane to that of assertability.

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." Now, it doesn't happen very frequently, but, as we all should suspect, newspapers do misreport scores from time to time. On several occasions, my paper has 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.

I'm interested here in accounting for the strange difference in assertability and knowledge between the lottery case and the newspaper case. Why does it seem you <u>do</u> have assertability and knowledge in and only in the newspaper case, where the probability of your being wrong is <u>higher</u>? I'll sidestep questions about whether you really do know in our two cases by focussing on explaining why it (at least) <u>seems</u> to us that you do know in the newspaper, but not in the lottery, case. Explaining this difference in the appearance of knowledge is, I think, an important step in any attempt to address the question of where we really do know, and is, in any case, puzzle enough for now.

A point of clarification. If, in the newspaper case, one were confronted by a skeptic determined to make heavy whether over the possibility that the paper's 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. But what I want to explain is why, <u>with no such a skeptic in sight</u>, we typically do judge that we know in the newspaper, but not in the lottery, case. (Unless so judging in the lottery case makes <u>us</u> skeptics, in which case I want to know why we're so naturally skeptics in the lottery, but not in the newspaper, case.)

II. The Solution

Although several candidate explanations suggest themselves quite naturally, the account I'll defend -- The Subjunctive Conditionals Account (SCA) -- is not one that immediately jumps to mind. According to SCA, the reason we judge that you don't know you've lost the lottery is that (a) Although you believe you're a loser, we realize that you'd 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'd 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 in [4] and is the explanation that would be suggested by Robert Nozick's theory of knowledge in [10]. But one needn't buy into Dretske's or Nozick's analysis of knowledge to accept SCA. (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 here by applying SCA to a variety of lottery- and newspaper-like cases and arguing that it outperforms its rivals. If I succeed in showing that SCA is the best explanation for why we have the intuitions we have, that should motivate us to seek an account of knowledge that makes sense of SCA without doing the violence to closure that Dretske's and Nozick's analyses do.[2,3]

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's (b)'s initial plausibility. If it can be shown to us that a subject would believe something even if it were false, that intuitively seems a pretty compelling ground for judging that the subject doesn't know the thing in question.

Second, there's this. In the lottery situation, even the most minute chances of error seem to rob us of knowledge and of assertability. It's puzzling, then, 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 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: 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.

But perhaps there's another explanation to be had.

III. No Determinate Winner, Losers

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's a loser. So you can't know your ticket's 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.

But this can't be the explanation. Even if the winner's already been picked in the lottery, so 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 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's 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. (If you insist that there is no winning ticket until it's been announced -- that it becomes a winner only at the announcement, not when the number's drawn -- then alter the case so that the winner _has_ been announced, but the people talking, though they know the announcement's been made, haven't yet heard what the winning number is.)

IV. The Existence of an Actual Winner: No-Winner Lotteries

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's a loser, there's the actual existence of a winning ticket, which is in relevant ways just like yours.[4,5] ("Somebody's gonna win.") By contrast, in the newspaper case, while there admittedly is a chance that your paper's 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. Perhaps only other copies of the edition I'm looking at are in the relevant ways like my copy. 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, supposing that all the tickets for the present drawing are alike in the relevant ways.

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 lottos, for example, usually have no winner. Still, it seems, you don't know you've lost. In case you think that's because the jackpot is carried over to the next month's drawing, 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 billionaire holds a one-time lottery, and you are one of the 1 million people who have received a numbered ticket. A number has been drawn at random from among 100 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 100 million; the chances that somebody or other has won are 1 in 100. In all likelihood, then, there is no winner. You certainly don't believe there's an actual winner. Do you know you're a loser? Can you flat-out assert you're 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.[6]

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

To approach this issue from the other side, what happens if we know that there actually is a "loser" newspaper? Suppose your newspaper announces that it's instituted a new procedure for checking and printing the results of 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's 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 is guaranteed to 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 said. "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 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 you know they've won. Indeed, suppose that in this new case you're asked whether you know if the Bulls won. I respond positively, as I'm sure almost anybody would.

Of course, this appearance of knowledge may fade in the presence of a skeptic determined to make heavy weather over the possibility that your paper's 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 I, our issue 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 in this and the previous section.

VI. 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 a puzzle of its own.

Here's the puzzle. If one's 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 clear 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 the conclusions to be drawn 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 so judge ourselves ignorant of that fact. And that's just what SCA predicts, since 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's only the belief in (b) that SCA predicts we'll be blocked from thinking is 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 won't be able to pay off the loan by the end of the year. Of course, I will 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 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 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's 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 millionaire whom I don't know at all will pick my name out of the phone book this year as her sole legal heir right before dying.

So things look pretty good for SCA. It's our beliefs in (b) and (b') that it predicts we won't think are 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 his skeptical urgings.[7] And, indeed, it is difficult -- in fact, intuitively repugnant -- 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 that former but not the latter. But these are all matters relevant to the issue of whether we really know. As for our particular judgments that we would ordinarily make as to what we know and don't know, and as for what we'd typically be willing to flat-out assert if asked, SCA gets things right.

VII. What about "My Paper's Accurate"?

You <u>believe</u> that your newspaper is accurate in the newspaper lottery case. But in that case, do you seem to <u>know</u>, and can you assert, that your paper's accurate when it comes to the sports scores it reports? Here we flip-flop. In settings in which we're focussed 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's a "loser" copy but in which we're not particularly focussed on that fact, we may judge that we do know. Being a big sports fan, I subscribe to the paper I do, as opposed to the other rag I could have ordered, after all, partly because I know its reporting of sports scores is almost always correct, as opposed to the competition, which frequently messes up. I've been looking at scores from this paper for a long time. If it were inaccurate, I'd have known it. This would all remain true if my paper switched to 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 focussed 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 you 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's 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's 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're a big sports fan who often looks at the scores and would have noticed if they were frequently wrong. If you're not thus like me, you may not seem to know your paper's accurate when it comes to its sports scores.)

VIII. That There's a Chance of Winning is the Whole Point of the Lottery!

We're now in a position effectively to put to rest a certain kind of proposed explanation that some readers may have been inclined toward ever since our puzzle was presented in section I. I won't spell this explanation out in full; I suspect that it can be completed in various significantly different ways. But the explanations I have in mind are all based on the observation that, with respect to the belief that one's lost the lottery, the chance that this belief is wrong -- i.e., the chance that you're a winner -- is intimately connected to the whole point of entering the lottery. By contrast, in the newspaper case, while there's a chance that your belief that the Bulls won is mistaken, this chance is not similarly connected to any of your goals.

The newspaper lottery in section V 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's 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 VI should show us why such a suggestion cannot provide the explanation we've been seeking. It's 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're wrong (i.e., the chance that your copy is the "loser") does prevent you from knowing, despite both its minuteness and its lack of a connection with any relevant goals.

IX. The Big Pay-Off, Etc.

Closely related to the proposed explanation explored in the last section is this slightly different, but equally doomed, proposal. Some suggest that it's 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 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're 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 you're not being the winner in the lottery case, where your interest in whether you're 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.

X. Probabilistic Thoughts and Statistical Reasons

Addressing 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. [5, p. 205]

Dudman doesn't identify the probabilistic underlying thought involved in the lottery case, but, presumably, it's something like this: Only one ticket out of the 20 million is a winner; so, probably my ticket's a loser. By contrast, in either of our newspaper cases (the regular 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's a loser. But 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 to 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 there are no relevant alternatives to P.[8] According to Cohen, while, in the cases in which we do think we know, there are alternatives to what we think we know, 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's 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 <u>relevant</u> 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. [1, p. 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 the lottery case because there our "underlying thought is reduced to relying on 'mere' probability." Could either of these be the explanation?

Well, first, 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/reasons that blocked knowledge and assertion in the lottery case, we should want to know <u>why</u> we think merely probabilistic thoughts only in the lottery case, where the chances that we're wrong are <u>less</u>.

But, more importantly, we don't just <u>happen</u> to think probabilistic thoughts in the lottery case: It seems they're forced upon us. Perhaps all we think in the newspaper case is the non-statistical and non-probabilistic: The newspaper says it; so, it's so. But why wouldn't similar, 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, <u>pace</u> Dudman, can survive probabilistic thought. Consider the original newspaper case. Those few incidents in which my paper's 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 very high; so they probably won; it's overwhelmingly likely." That's 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. If nothing's been said in our conversation about those rare instances of mistakes, 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's 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.

XI. Admitted Possibilities of Error and Admitted Probabilistic Grounds

But what if the possibility of error, the instances of past errors, and the probabilistic grounds are not just privately thought, but are publicly interjected into the conversation? Would <u>that</u> preclude assertion and block the appearance of knowledge?

Still addressing the issue of assertability in the lottery, and immediately after the passage quoted above at the start of section IX, Dudman continues:

> Since mere probability leaves open the contrary possibility, assenting on the basis of admitted mere probability would be to treat admitted possibilities of to the contrary as if they were impossibilities. This is apparently why we steadfastly withhold assent in the face of rocketing odds. [5, p. 205]

We've now moved beyond underlying thoughts. Here it's <u>admitted</u> merely probabilistic grounds and <u>admitted</u> possibilities to the contrary that are fingered as the assertion killers in the lottery. But this can't be right. In what sense have we (must we have) <u>admitted</u> the possibility that we've won the lottery? The possibility invariably seems to be there, staring us in the face and blocking the assertion that we've lost, whether we do anything at all to admit it's there or not. But, even though one seems to be present, no such possibility seems inevitably to block assertion in the newspaper case, so long as we keep quiet about the possibility.

But suppose that such things as probabilistic grounds, instances of past errors, and possibilities to the contrary aren't kept quiet, but are admitted out loud in the newspaper case, for instance, as follows:

A: How do you always know so many basketball results?

B: I check my paper every morning before I come to work.

A: Papers make mistakes sometimes, you know.

B: I know. How often do they mess up?

A: Oh, <u>very</u> infrequently. I just wanted you to know that it's always <u>possible</u> that they're wrong. Your grounds are probabilistic.

B: I realize that.

A: But don't worry. The probability that your paper's right in any given case is very, very high. The likelihood that it's wrong is so low as to be not worth considering. So, did the Bulls win last night?

If you're B, and your only source of information for the result of the Bulls game is your newspaper's report that they won, which is presently before your eyes, how do (can) you respond here? The rambling,

> They probably won. This paper reports them as winning, and, as you just said, the probability that it's correct in any given case is very high,

is perhaps admissible, though it's not what I'd say. Seemingly preferable and certainly allowable is the simple, "Yes, they won" (perhaps followed by, "Says so right here"), which is how I would respond. Assertability here survives the admitted possibility of error and an admittedly probabilistic basis.

So does knowledge. In the above dialogue, change A's closing question to: "By the way, do you know if the Bulls won last night?" I answer, "Yes, I do. They beat the Knicks 95-83." By contrast, if A and B are discussing how very low the chances are that any given ticket has won the Lotto drawing ("So low as to be not worth considering," A has said), B cannot assert he's lost, and doesn't seem to know it.

So, in the newspaper case, assertability can survive probabilistic and statistical thoughts and reasons, even where these are publicly admitted. And in the lottery case, we seem prohibited from relying on the simple, non-probabilistic: "It's a Super Lotto ticket (for heaven's sake!); so, it's a loser," to make the unqualified assertion "I've lost," even if we haven't uttered a probabilistic or statistical word. So the suggestions we've looked at in this and the previous section can't explain the divergence in assertability and in apparent knowledge between our cases. Again, SCA can.[9]

REFERENCES

1. Stewart Cohen, "How to Be a Fallibilist", Philosophical Perspectives 2 (1988): 91-123.

2. Keith DeRose, "Epistemic Possibilities", Philosophical Review 100 (1991): 581-605.

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

4. Fred Dretske, "Conclusive Reasons", Australasian Journal of Philosophy 49 (1971): 1-22.

5. V.H. Dudman, "Probability and Assertion", Analysis 52 (1992): 204-211.

6. Gilbert Harman, "Knowledge, Inference, and Explanation", American Philosophical Quarterly 5 (1968): 164-173.

7. David Lewis, "Elusive Knowledge", Australasian Journal of Philosophy 74 (1996): 549-567.

8. G.E. Moore, Commonplace Book: 1919-1953 (London: Allen and Unwin, 1962).

9. Robert Nozick, Philosophical Explanation (Cambridge, MA: Harvard UP, 1981).

10. Peter Unger, Ignorance: A Case for Scepticism (Oxford: Clarendon Press, 1975).

NOTES

1. For an impressive defense of (b) by Peter Unger, see [11, pp. 250-271]. In [2], I give further support to (b) by using it to explain various linguistic phenomena (pp. 596-605). Cf. G.E. Moore's treatment of "Dogs bark, but I don't know that they do" [9, p. 277].

2. In [3], I present a contextualist account of knowledge attributions that does just that.

3. SCA is also closely related to the explanation Gilbert Harman proposed to our puzzle in [6]. (Harman seeks to differentiate the lottery case from what he calls the "testimony case," where the latter is a situation in which "a person comes to know something when he is told about it by an eyewitness or when he reads about it in the newspaper" [6, p. 166].) His explanation is that in the newspaper case, one knows because one can there perform an inference to the best explanation from one's evidence to the fact that the Bulls won. By contrast, no such inference to the best explanation is available to the conclusion that one will lose the lottery (though one can infer, on statistical grounds, that one will probably lose). The causal connection between the Bulls winning and one's evidence that they've won, which makes possible this inference to the best explanation, also grounds the truth of the subjunctive conditional crucial to SCA: If they hadn't won, I wouldn't believe they had. SCA does sometimes produce different results from Harman's account, though. For instance, Harman's view would suggest that I can't know that my neighbor's copy of the paper reports the Bulls as winning, based on the fact that my copy of the paper reports them as winning, since the fact that my neighbor's copy lists the Bulls as winning doesn't explain how my copy of the paper came to report them as winning. By contrast, SCA doesn't predict that I won't seem to know here. If my neighbor's paper hadn't reported the Bulls

as winning, that would probably be because they didn't win. But then my paper wouldn't have reported them as winning, and I wouldn't now believe that my neighbor's paper reports them as winning. This case favors SCA over Harman's account, I think, because we would ordinarily judge that I can know that my neighbor's paper lists the Bulls as winning. Suppose that, as my wife and I watch that Knicks fan pick up his paper, I say, having read my paper, "He'll be disappointed when he reads that." When my wife asks me why, I can flat-out assert, "Because it says the Bulls won last night." Still, I consider SCA to be a close relative of Harman's account. There are, of course, ways of making Harman's account give the right result here. These, I think, will move his account closer to SCA.

Chapter 8 of Harman's later [7] modifies the account of knowledge in [6]. But in [7] Harman's treatment of the lottery problem involves a denial that one even believes that one will lose the lottery. I think this is a mistake, but I don't, in the present paper, address solutions according to which one doesn't even believe one will lose the lottery.

4. In cases in which the winning ticket hasn't yet been picked, the claim here will be that there will be an actual winning ticket just like yours.

5. The solution David Lewis advances in [8], above (see pp. xxx, xxx), based on his rules of Resemblance and of Actuality, is an explanation of this rough type. See note 6, below.

6. These no-winner cases (the state lotto with no winner, and the billionaire's lottery), I think, can be handled by Lewis's account (see the reference in note 5, above). For while there are no winning tickets in these cases, there are (or will be) winning numbers, and the possibility that your number will be chosen saliently resembles actuality (that some other number will be

drawn). But Lewis's account runs into trouble with other no-winner cases, like the following variant of the billionaire's lottery. Suppose there are only one million numbered balls, one corresponding to each of the one million tickets held by players of the lottery. But there's a .99 probability that the random device set up won't pick up any numbered ball. If a numbered ball is picked up, of course, the holder of the corresponding ticket receives the prize; if, as is likely, no ball is picked up, nobody wins. Here, I still don't seem to know I'll lose, though the possibility that I believe obtains -- that nobody wins -- doesn't saliently resemble the possibility that I win: The former involves no ball being picked, and is <u>far</u> more probable than is the latter.

7. In such skeptical arguments, not-(b) and not-(b') function as <u>skeptical hypotheses</u>. We then have instances of the time-honored pattern of skeptical reasoning which proceeds from the intuitive claim that we don't know some skeptical hypothesis to be false to the conclusion that we don't know some other thing (here (a) or (a')), which we ordinarily do take ourselves to know. In [3] I address such skeptical arguments in a way that utilizes SCA.

8. According to most versions of the relevant alternatives theory, S knows that P iff (roughly) S has a true belief that P <u>and can rule out</u> 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 (see p. 101). Thus, in cases in which S does know, what most versions of the relevant alternatives theory classify as relevant but ruled out alternatives are for Cohen irrelevant alternatives.

9. Thanks to Stewart Cohen and Christopher Hitchcock for helpful comments.