DeRose on Lotteries

Peter Baumann

This is a pre-print of an article published in the *International Journal for the Study of Skepticism* 10, 2020, 44-67. The final authenticated version is available online at: https://brill.com/view/journals/skep/10/1/article-p44_44.xml [https://doi.org/10.1163/22105700-20191369]

Abstract

This article discusses Keith DeRose's discussion of the lottery problem in chapter 5 of his recent *The Appearance of Ignorance*. I agree with a lot of it but also raise some critical points and questions and make some friendly proposals. I discuss different ways to set up the problem, go into the difference (quite relevant here) between knowing and ending inquiry, propose to distinguish between two different kinds of lotteries, add to the defense of the idea that one can know lottery propositions, give a critical discussion of DeRose's contextualist solution to the problem, and support his defense against an absurdity objection with additional arguments.

Keywords

Knowledge; lottery problem; skepticism; ending inquiry; epistemic lotteries; DeRose

1. Introduction

Keith DeRose's impressive recent *The Appearance of Ignorance* has much to say about many topics—and more than can be dealt with at article length. I will focus here on DeRose's very own analysis of and solution to the lottery problem. It so happens that I agree with a lot of what DeRose is saying about lotteries. So, my comments will be a mix of friendly proposals, questions and criticisms. I thus cannot offer the drama of deep disagreement here; however, a more detailed discussion of a complex view like DeRose's can be worth a lot, too. It will be useful to start with an exposition of the lottery problem (which differs a bit from DeRose's exposition of it) (Section 2). Given this background, I will then (Section 3) turn to DeRose's somewhat different way of looking at the problem, using the contrast between reading (in a newspaper) about a lottery outcome and reasoning probabilistically about the outcome. Section 4 continues this by focusing on the relation between knowing and ending inquiry and raises doubts about one presumed difference between reading and reasoning about lottery outcomes. Section 5 proposes a distinction between two kinds of lotteries which should be useful here. Section 6 defends, supporting DeRose, the controversial claim that one can know lottery propositions. Section 7 gives a critical discussion of DeRose's contextualist solution to the lottery problem. Section 8 gets back to Section 6 and proposes further replies (further to DeRose's) to one particularly serious objection against the claim that one can know lottery propositions.

2. The Lottery Problem

Suppose Jackie owns a ticket in a fair lottery with lots of tickets (say 1 Million) and exactly one winning ticket. She hasn't heard yet about the results of the drawing but she is aware of the poor chance of her winning. She has no other evidence relevant to this case and comes to believe that her ticket is a losing ticket. Suppose finally that this is true. Can she know—even before she finds out about the drawing and just on the basis of the probabilistic evidence—that her ticket is a losing ticket?[1]

[1] There could be a winning ticket but no winner because nobody has managed to buy the winning ticket. If there isn't even a winning ticket or if there isn't even a losing ticket, then it's not a lottery in the sense relevant here any more. If there is more than one winning ticket or if the chances of winning differ for different tickets, nothing much about our case changes. Even if your ticket has a 0 (100) percent chance of winning but you are not aware of this, then one might still consider it plausible to claim that you don't know that you're losing (winning). In this case, the probabilistic justification for your belief would be based on the false assumption that there is a non-zero chance of winning as well as of losing, which

Current orthodoxy answers in the negative. If that is the correct answer, then we have another type of a Gettier case: a justified true belief which does not amount to knowledge (Gettier 1963). However, not everyone agrees that this negative verdict is true or even obviously true (as many nowadays seem to assume).[2] But let us assume for now that a subject in Jackie's kind of situation does not and cannot even know that their ticket is a losing one.

Suppose further that Jackie—who is well aware of her somewhat dire financial situation—believes that she will not be able to afford a much bigger apartment any time soon. Assume that this is true. Can she know that she won't be able to afford a much bigger apartment any time soon?

Current orthodoxy answers questions like this one in the positive (even overwhelmingly so). Everyone who isn't a skeptic about knowledge anyway would agree. The verdict that Jackie can know this is very plausible, to say the least.

Finally, it seems very plausible to claim that if Jackie knows that she won't be able to afford a much bigger apartment any time soon, then she knows that she won't win this lottery. Equivalently, if she doesn't know the latter, then she doesn't know the former; one cannot know the one but not the other.

---

assumption, in turn, would be based on the true assumption that "safe-ticket"-lotteries (with no losers or no winners) are very rare. However, I am more inclined to say that anything including safe tickets (or also "rigged lotteries") are not lotteries in the sense relevant here. Here, we can focus on "normal" and fair lotteries with exactly one winner and many losers. (For the case in which someone owns all the tickets but one, see below.)

[2] Keith DeRose (2017: 133 n.3) denies that this is obvious. He also denies that it is true *simpliciter* (see below).

Together, our three claims constitute an inconsistent triad and also a problem and puzzle because all three of them are quite plausible:

Jackie's Puzzle

(i) Jackie does not know that (a) she won't win the lottery;

(ii) Jackie does know that (b) she won't be able to afford a much bigger apartment any time soon;

(iii) If Jackie does not know (a), then she does not know (b).[3]

One way to support (iii) is extremely straightforward. Jackie's knowledge that she won't be able to afford a much bigger apartment any time soon should be described more precisely not as knowledge of the proposition *I won't be able to afford a much bigger apartment any time soon* but rather of a more complex proposition like *I won't be able to afford a much bigger apartment any time soon because I won't inherit a lot of money soon and I won't find a suitcase full of money in the attic and I won't win some lottery*. If knowing a conjunction *ipso facto* involves knowing each conjunct, then by knowing the fully specified (b) Jackie knows that she won't win the lottery. This gives us the following:

Jackie's Strengthened Puzzle

(i) Jackie does not know that (a) she won't win the lottery;

(ii') Jackie does know that (b') she won't win the lottery;

(iii') If Jackie does not know (a), then she does not know (b').

[3] We can even strengthen "does not know" to "cannot know".

Here, the first two claims on their own create a problem: Jackie does and does not know that (a)/(b'). Such inconsistency would be most unwelcome.[4]

Currently, the major strategy for supporting (iii)—if it is not accepted as intuitively plausible—consists in invoking some principle of closure of knowledge under known entailment.[5] Here is simple version of it:

(Closure) If S knows that *p*, and if S knows that *p* entails *q*, then S knows that *q*.[6]

---

[4] If, alternatively, we were to claim that knowing a conjunction does not involve knowledge of the conjuncts but only allows one to know the conjuncts (by inference), then we would get to the contradiction more indirectly and with a support of a principle like the closure principle. On that now.

[5] There has been some discussion recently about the question whether a closure principle is required for supporting (iii) or whether there are other principles, e.g., a principle of underdetermination that could do the same. It is not necessary to go into this here.

[6] It should not surprise that things aren't quite that straightforward. Closer to the truth is a principle like the following:

> If S knows that *p*, and if S competently infers *q* from *p,* and thereby comes to believe that *q*, then S knows that *q*. (See also DeRose 2017: 164–165; see also Baumann 2016 and 2011).

I am pessimistic about the prospects of identifying sufficient conditions for inferential knowledge from knowledge (and of finite chisholming; see DeRose 2017: 164). Fortunately, nothing much depends on the details here and at least for now we can work with (Closure). We also don't need to go into the details of the nature of entailment here or into the most precise specification of propositions. Principles like

If we can assume that the Jackie knows the entailment *If I won't be able to afford a much bigger apartment any time soon, then I won't win the lottery*, then, given (Closure), we can infer *If Jackie knows that she won't be able to afford a much bigger apartment any time soon, then she knows that she won't win the lottery*—which is, of course, equivalent to (iii).

As can be seen easily, Jackie's puzzle generalizes to all subjects. It also generalizes to all pairs of propositions (*l*, *o*) which are such that one (*o*) is an ordinary proposition we typically (if not in a radical skeptical mood) take ourselves to know while the other one (*l*) is a so-called "lottery proposition". Surprisingly, there have not been many attempts at all to give a more precise explanation of the term "lottery proposition". But we can say that not just propositions about "literal" lotteries are lottery propositions. It is not a matter of content but of the subject's epistemic relation to the proposition. For instance, one can create the above problem with this pair of propositions (for cases like this, see, e.g., Hawthorne 2004): *I know that my car is in the garage* and *I don't know that my car hasn't been stolen recently (which very rarely happens) and relocated elsewhere*. What would make a proposition a lottery proposition depends on the nature of the evidence we have for it rather than on its content: merely probabilistic or not.[7] This

---

(Closure) also give consolation to the ignorant or to lazy thinkers: if one doesn't know the entailment or doesn't make the relevant inference, then one can know the *p* without knowing the *q*. Not knowing one thing (the connection between *p* and *q*) might enable one to know another thing—which one would not know if one knew the first thing. This is very interesting but I won't go into it here.

[7] If all or almost all or even many propositions have merely probabilistic support, then our problem would collapse if not more could be said about the nature of lottery propositions. Hence, more should be said, at least in that case. Here, I will rely on a to a large degree still intuitive grasp of the nature of lottery propositions.

implies that one and the same proposition (that my ticket has lost) can be a lottery proposition for me now (not having heard about the drawing yet) while not for you later (having then read the papers). So, we can generalize from *Jackie's Puzzle*[8]:

Lottery Problem

(1) S does not know that *l*;

(2) S does know that *o*;

(3) If S does not know that *l*, then S does not know that *o*.

Since a lot (if not all) of ordinary propositions entail some lottery proposition, a very general problem arises. How can we know any of the things we ordinarily assume we know if (1) and (3) hold? This is a serious skeptical problem, too, and it very much parallels the skeptical problem nowadays discussed as the problem of Cartesian Skepticism (see DeRose 2017: 153). It does so without involving any farfetched or outlandish scenarios (which is an advantage for the skeptic). As a template for Cartesian Skepticism we can use what DeRose calls the "Argument from Ignorance" (with "H" for some skeptical hypothesis like the one that I could be a brain in a vat, and with "O" for some ordinary proposition entailing not-H):

Skeptical Problem

"1. I don't know that not-H.

2. If I don't know that not-H, then I don't know that O.

[8] Each of the three claims holds with necessity. For simplicity's sake and because nothing depends on it here, I will not explicitly mention this in the following.

So, C. I don't know that O." (DeRose 1995: 1)

The lottery problem allows for a range of responses.[9] One could deny (3) or what supports it, namely something like (Closure), for instance. Not many would like to do that. Or one could deny (2) and be a skeptic about ordinary knowledge. This is, to most, not a very attractive option. In contrast, one could be a "Moorean" about lottery cases and deny (1). Some but not that many would choose this strategy. Finally, one could choose a contextualist option according to which the truth conditions for (1) and (2) vary with context of use (and in such a way that either both *l* and *o* are known or that both are not known).[10] DeRose is not only one of the most prominent defenders of the last option,[11] but also offers a quite unique solution.[12]

[9] For the problem itself, see, apart from DeRose (1996) and amongst many, Harman (1968), Harman (1973: 160–161), Harman & Sherman (2004), Vogel (1990: esp. 16, 22), and Hawthorne (2004). Further interesting recent discussions of the problem are to be found in Olen (1977: 521–523), Stemmer (1982) Adler (1986: 244–248), Dudman (1992: 205), Ryan (1996: 130), Nelkin (2000: 388–390), Weintraub (2001), Olin (2003: 98–104), Greco (2003; 2004), Roush (2005: 132–133). See also Brueckner (2012: sec. V) as well as the reply in Roush (2012: sec. V), Collins (2006), Douven (2007) Nagel (2011), Mills (2012), Bondy (2013), McKinnon (2013), and Timmerman (2013). For experimental data, see Turri & Friedman (2014) and Ebert, Smith & Durbach (2018). The lottery problem should be clearly distinguished from Kyburg's different lottery puzzle (see Kyburg 1961: 197–198). I won't discuss any of the literature mentioned here in more detail; I will rather focus pretty much exclusively on DeRose.

[10] This is by far not a trivial assumption. Why couldn't there be some context such that *o* is still known but *l* isn't? I will leave this issue aside here.

[11] For different versions see, e.g., Cohen (1988) and Lewis (1996).

3. Newspapers and Lotteries

It is worth stressing that in chapter 5 of *The Appearance of Ignorance* (as before in his 1996) DeRose gives a different exposition of the lottery problem than the one above. He is not talking about a different problem but he is focusing on different aspects of the same complex problem. I wonder why that is and why he, of all people, doesn't focus more on those aspects of the lottery problem that are very much parallel to the skeptical problem. Why does he focus so much on the other aspects of the lottery problem? Anyway, let us look at his perspective on the issues.

DeRose starts with three observations. First, it is intuitively plausible that a given lottery ticket (in a normal lottery: see above) is a loser (2017: 132). The probabilities of losing are just very high.[13] Second, it is intuitively plausible that one knows what one read in the newspaper, namely, for instance, that the Bulls won their basketball game, and this is so even though the probability that the newspaper got it wrong is much higher than the probability that one's ticket is a winner (132–133). Finally, it is also intuitively plausible that if one doesn't have knowledge that one's ticket will lose, then one also doesn't have knowledge that the Bulls won (133–134). The mutual incompatibility of these three claims creates a puzzle and problem:

---

[12] It seems to me that amongst authors writing about the skeptical problem and the lottery problem there is clearly more doubt about the possibility of knowing lottery propositions than about the possibility of knowing the falsehood of some skeptical scenario. Given the parallel between the skeptical problem and the lottery problem, this is surprising and in need of explanation.

[13] One should add here explicitly that this is all the evidence the subject has (they haven't heard the news about the drawing yet, etc.). DeRose seems to assume this implicitly.

(4) Keith does not know that this ticket is a loser;

(5) Keith knows that the Bulls won;

(6) If Keith does not know that this ticket is a loser, then he does not know that the Bulls won.

Why did DeRose choose *The Bulls won* and not *This ticket is a loser*? To be sure, in the latter case the third claim would be a trivial logical truth (If he doesn't know, then he doesn't know). But one could, alternatively, have explained the puzzle without using (6): how come that Keith can know (after reading the paper) that the Bulls won (or that his ticket lost) though he doesn't know that this ticket is a loser based on probabilistic reasoning—given that the chances of error are much greater in the first case? Apart from that, it will turn out later that it would be good to choose the same proposition both in the case where the subject merely considers probabilities and in the case where they're reading the papers.

DeRose proceeds in two steps. First, he proposes a mere explanation of why we think differently about the regular lottery[14] and the newspaper case and how that would make sense of the idea that the subject's epistemic position is stronger in the newspaper case than in the regular lottery case (sections 5.1–5.13). In the second half of the chapter (sections 5.14–5.26), he proposes a contextualist justification and theory about all this.

First, the explanation. DeRose proposes his well-known Subjective Conditionals Account (SCA) here. The belief that this ticket is a loser is insensitive: the subject would believe this even if the ticket won. The belief that the Bulls won is, in contrast, sensitive: the subject would not

[14] From now on I'll often call this "regular lottery" because later on a different kind of lottery will play an important role.

believe this if the Bulls had lost. We tend to judge that a subject does not have knowledge if their belief is insensitive (but we tend to judge that the subject might have knowledge if their belief is sensitive) (2017: 137–138). DeRose is careful to stress that he does not want to claim that sensitivity (of true belief) is a necessary or sufficient or necessary and sufficient condition of knowledge (137). Again, at this point he is not proposing a view about knowledge but rather trying to explain the phenomena.

There are several questions about this. First, why choose sensitivity as an explanation? Couldn't, for instance, the idea of safety also do the job? Not easily (one could argue) could the belief that the Bulls won have been wrong but easily could the belief that this ticket is a loser have been wrong. And couldn't one suggest, at least, that we tend to judge that a subject has knowledge only if their belief is safe? Or couldn't a relevant alternatives account do the same job and not a worse one than sensitivity? That one's ticket is a winner is a relevant alternative while that the Bulls lost isn't? I don't want to defend any of these alternatives here but rather point out that it's not that obvious that sensitivity has an advantage here. To be sure, DeRose considers a good number of alternative explanations and criticizes them, often very convincingly (sections 5.3–5.13). However, one doesn't have to be a friend of safety accounts (or relevant alternative accounts) of knowledge to wonder why safety (or the idea of relevant alternatives), for instance, is being left aside. Interestingly, later on (2017: 154 n.10) he briefly mentions safety accounts as potentially helpful ones. It is known (see here 154 n. 10) that DeRose is open about alternative modal or non-modal accounts of knowledge, but since sensitivity plays a major role in his contextualist account of knowledge (see section 5.14 or his 1995), more of a defense of sensitivity would be useful here.

Apart from this, one wonders about the details of the idea of sensitivity. It is very common to put it in terms of possible worlds and a closeness metric between them. Here is one way to do this.[15] In close possible worlds in which this ticket wins, the subject still believes that it is a loser; in this sense their belief is insensitive. In contrast, in close possible worlds in which the Bulls lost, the subject does not believe that the Bulls won (because in those worlds the newspaper reports correctly); in this sense their belief is sensitive. Surprisingly, DeRose later (see 154 n.10) briefly remarks that he does not want to use the idea of a range of possible worlds here because of problems like the one that one would have to consider winning the lottery a distant possibility. (Why that? Why not consider it a close possibility?) He doesn't say much at all here about these problems with such a possible worlds semantics. To be sure, there are serious problems with the idea of a closeness metric between worlds. But one wonders how else one could make sense of the notion of sensitivity? (DeRose also briefly mentions (in the same footnote) that perhaps safety accounts can do better here but this is not easy to see, given that they use the idea of possible worlds and a closeness metric in the same way sensitivity accounts do.) Finally, he makes another brief but very interesting remark (see 140): what if there are some copies of the same newspaper which report the outcome of the game (or of the lottery) incorrectly? DeRose finds it intuitively plausible to say that such misleading copies are not in all relevant ways like the (correctly reporting) copy the subject is reading. What is the relevant dissimilarity? Is it the incorrectness of the report? Is it that the subject knows that the Bulls won because their copy reports correctly and because we can put aside the possibility of misinformation? However, the chance of error seems utterly relevant to the possibility of knowledge.[16]

[15] Not more than a rough version is needed here.

[16] We will get back to the possibility of incorrect newspaper reports below.

4. Knowledge and Inquiry

The difference between the regular lottery case and the newspaper case is interesting enough to deserve some more comments. Consider two subjects, Jackie and Jo, who share a lottery ticket (as a matter of fact, a loser). At the beginning, neither subject has any evidence concerning losing or winning apart from the probabilities. Let us, for the moment, suspend judgment about whether at that time they know or don't (can't, even) know that their ticket is a loser. Suppose finally that Jo but not Jackie later reads in her newspaper that their ticket has lost. That is the moment when Jo's epistemic position regarding *Our ticket has lost* improves and becomes better than Jackie's epistemic position regarding *Our ticket has lost*. According to the intuitive judgments DeRose is talking about, this difference in epistemic positions is one between knowledge and the lack thereof. But, again, let us assume that this is an open question and try to get closer to an answer by raising a related question: what good reason might someone have to read the newspaper report about the lottery outcome?

Would Jo, after reading the paper, say to Jackie "Now we know: we've lost!"? There is doubt that this would be the normal and appropriate thing to say. Much more appropriate and common would, I think, seem something like "No surprise here: we've lost!".[17] And then the case would be epistemically closed[18] for Jo and Jackie. If we can make the (hopefully plausible) assumption that

[17] Perhaps even "We've lost: I knew it!". But this is controversial here and also raises complicated questions about the first person use of "know".

[18] I am using terms like "epistemic closing" in a normative sense: not continuing inquiry only counts as closing if there are good reasons for not continuing it.

(CK) An epistemic case isn't epistemically closed for a subject if the subject doesn't yet have the relevant knowledge,

then Jackie and Jo have knowledge that they've lost at least by the time Jo has read the paper. And, if what seems appropriate and inappropriate to say after reading the paper—not "Now we know!" but something else (see above)—is indeed what is the right or wrong thing to say, then they already have knowledge that they've lost before they've read the paper. The reason to read the paper would then rather be that even though knowledge is necessary for epistemically closing an epistemic case, it is not also sufficient for it. What Jackie and Jo want in addition to knowledge might be a maximum or a very high degree of certainty (out of theoretical interest like curiosity or out of practical interest, especially when the prize is non-trivial) or a significant increase in the degree of certainty—which they get because of the independence of the new evidence.[19]

Consider now a third subject, Fortunata (a neighbor), who also owns a ticket in the same lottery—but this time a winner. Like Jackie and Jo, she first only has probabilistic evidence favoring a loss. Then she reads in the paper that she has won the lottery. If Fortunata "can't believe it", then a normal and appropriate thing for her to say then would not be "I've won. I know it: the Paper says so!" but rather something like "I can't believe it: I've won. Look at this, the paper says my number is a winner. Am I dreaming? Could this be true? I need to check and make sure this is correct!" Fortunata doesn't come to know, by reading the paper, that she has won. But this is not because one cannot come to know such things by reading the paper, but rather because the surprising and unlikely nature of the news makes it reasonable (both for

[19] In addition, in all such cases expected utility might also give reasons to check further.

theoretical or epistemic reasons which suggest to treat very surprising evidence with caution, and for practical reasons, given the risk of being wrong after all) for her to suspend judgment a bit or at least lower her degree of conviction that she has won to a degree not sufficient for knowledge. (Knowledge that *p* requires a sufficient degree of conviction that *p*.)

If, however, Fortunata does believe what she has read—that she has won—then the normal and appropriate reaction wouldn't be to leave it at that and spend money celebrating but rather to look for more and independent evidence. Why would it be unreasonable for Fortunate to close the epistemic case upon reading? Is it because she doesn't have knowledge yet? If, apart from (CK), it would also be the case that

> (KC) A proposition isn't known by a subject if the epistemic case isn't epistemically closed for the subject yet,

then Fortunata wouldn't get to know that she has won by reading the newspaper while Jackie and Jo knew that they have lost, at least by the time they had read the papers. This asymmetry seems implausible, apart from the fact that it is hard to see why one couldn't come to know things like that by reading the paper (especially in cases of lotteries with only a small or no prize). It is better to reject (KC) while keeping (CK). Fortunata knows in this case that she has won the lottery (if she can believe it), but the case is not yet epistemically closed for her. It is reasonable for her not to close it yet and to look for more and independent evidence, again for both epistemic and practical reasons (to gain certainty, to minimize risk, etc.).

All this suggests that facts about whether one has reason to continue inquiry about whether *p* or rather to close the epistemic case about whether *p* don't have much to tell us about whether the subject has or lacks knowledge that *p*. I don't want all this to be understood as an argument in favor of the possibility of knowing lottery propositions (for that, see below), but rather as an attempt to defend that possibility against one particular objection: that the fact that one has a reason to read the papers about the lottery results shows that one didn't already know one's ticket has lost.

The above also suggests that there is "pragmatic encroachment" for the epistemic closing of epistemic cases (though it does not suggest this for knowledge): whether an epistemic case can be closed depends on the practical stakes. Let "End" stand for "no further checks after reading the newspaper" and "Cont" for "further checks after reading the newspaper:

	loss	win
End	A	B
Cont	C	D

In the case of a loss, End is the appropriate act. The greater the difference between A and B and the greater the probability of a win, the greater the stakes for End. In the case of a win, Cont is (plausibly) the appropriate act. The greater the difference between C and D and the greater the probability of a loss, the greater the stakes for Cont. The lower (higher) the stakes, the more reason the subject has (not) to close the case.[20] It would be very interesting to know what DeRose would say about all this.

[20] Jessica Brown's (2008: sec. 7) surgeon case could also be used to support this claim. On the notion of stakes, see Baumann (MS).

5. Newspaper Lotteries

In sections 5.7 and 5.8 of *The Appearance of Ignorance*, DeRose focuses explicitly on cases of newspapers that are overall very reliable but not completely: some of the copies of some editions contain incorrect reports about last night's Bulls game or about the last lottery results. DeRose uses the term "newspaper lottery" here (2017: 143).

It is worth taking a closer look at the term "lottery" in this context. Reading such a newspaper is a very reliable epistemic method of acquiring a true belief about some subject matter (about whether *p*); however, a few uses of the method lead to false beliefs. Like in a regular lottery, there are many items (uses of the method, corresponding to individual tickets) and an interesting property P shared only by a few of the items (resulting in a false belief, corresponding to being a winning ticket). One can thus characterize DeRose's "newspaper lottery" as an "epistemic lottery":

> (Epistemic Lottery) S plays an epistemic lottery just in case the epistemic method of belief acquisition used by S is not completely reliable.[21]

One may very well wonder how useful the notion of an epistemic lottery is, given that every method of belief acquisition (or almost every or at least many) is not completely reliable. However, and more importantly, the use of the term "lottery" both for regular lotteries and for epistemic lotteries, can occlude important differences. In DeRose's case of a reader who consults

[21] We can put aside here the question what constitutes reliability (something probabilistic, something modal, etc.).

the paper to find out whether the Bulls won, the subject (the reader) is playing an epistemic lottery but not a regular lottery. The subject's belief is about the outcome of a game, not about the outcome of a regular lottery. In contrast, a subject holding a ticket and trying to figure out whether the ticket is a loser, could either use probabilistic reasoning or consult the paper; in both cases they are playing an epistemic lottery. But in addition, they're also playing a regular lottery. They are trying to acquire a true belief about the outcome of a regular lottery (the belief's content being a lottery proposition). So, we have two levels of analysis here:

> Level I: the content of the relevant beliefs: a lottery proposition and about a regular lottery, and
>
> Level II: the playing of an epistemic lottery in order to acquire such beliefs.

Someone who is playing an epistemic lottery is not (at least not typically) trying to find out whether they have won the epistemic lottery; this would just start an epistemic regress. Rather, they are trying to find out something else (whether the Bulls won, what the lottery results are, whether it will rain later that day, etc.). It is, I think, crucial, to keep these two dimensions—or these two different kinds of lotteries—clearly distinguished.

DeRose continues by claiming that we would still judge that a reader who gets the correct result about the Bulls' game from the paper gets to know that the Bulls won. Similarly, a reader who's playing this particular epistemic lottery, the newspaper lottery, could find out that their ticket has lost. The above questions about why one should think so and why one should continue to think that mere probabilistic reasoning does not lead to knowledge are still open (but see below). DeRose only briefly (144) mentions the possibility that a skeptic might use the fact that there are some misleading copies around to deny knowledge for the reader. The question what exactly stands in the way to knowledge and what not is still open, and not just for some skeptics and epistemological troublemakers.[22]

DeRose begins section 5.8 with this question: “Isn’t the newspaper lottery case *just like* the regular lottery?” (145). Using the above distinction of two types of lotteries (which DeRose doesn’t), we can answer with “Yes and No!”. In both cases the subject is playing an epistemic lottery. But only in one case is the content of the relevant belief a lottery proposition and about a regular lottery (the subject is playing an epistemic lottery to find out about a regular lottery); in the other case it is not a lottery proposition but about the outcome of the Bulls’ game. Here it is useful to replace DeRose example proposition *The Bulls won* by *This ticket is a loser*. As far as the possibility of knowledge is concerned, it shouldn’t matter, one thinks, what exactly the reader learns from the paper. And in order to find out what determines knowledge we should reduce the number of irrelevant variables. If we then compare the case of probabilistic reasoning about the outcome of the regular lottery with the case of reading the paper in order to find out what the outcome of the regular lottery is, we will have to answer the above question with “Yes and Yes!”. In both cases, the subject is playing an epistemic lottery, and in both cases this epistemic lottery concerns a lottery proposition. This sharpens the puzzle more than DeRose’s presentation does.

DeRose doesn’t seem to accept the above distinction between epistemic lotteries and lottery propositions: “… 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

---

[22] DeRose (2017: 154 n. 10) does not want to say that, given that one has picked a correct copy of the paper, the possibility of having picked an incorrect one is remote while the possibility of holding a winning ticket, given that one holds a loser, is a close one.

'loser' newspaper" (145). This comparison is one between Level I for the regular lottery case and Level II for the newspaper case. This is a bit confusing. Why not compare the regular lottery case and the newspaper case across the same level? Why not point to the fact, again, that we would judge that the probabilistic reasoner doesn't come to know their ticket has lost while the reader does know that (and also that the Bulls won) (145)? DeRose thinks that SCA can at least explain this, but this question is also still open (but see below).[23]

DeRose also adds that as much as the reader knows that the Bulls won even though they don't know they've won the newspaper lottery (one could add: typically, they wouldn't even have a belief about this), so the probabilistic thinker doesn't know they've lost the lottery though they do know that they won't be able to pay their debt (145). This remark is puzzling for at least two reasons. First, the analogy seems misconstrued. In the newspaper case, DeRose contrasts knowledge at Level I with lack of knowledge at Level II; in the other case, he contrasts two propositions—one a lottery proposition and one presumably not—without mentioning Level II. And why bring in a second proposition? Second, if I know that I won't be able to pay back my debt, then (given closure and given that I know the entailment that if I won't be able to pay back my debt, then I have lost the lottery), then I know that I have lost the lottery. However, this leads us back to *Lottery Problem*. It is somewhat surprising that DeRose does not go further into this (again, see below). In any case, there seems to be more of a disanalogy than an analogy in

[23] It is interesting and worth mentioning here that DeRose, who coined the term "abominable conjunction" (see DeRose 1995: 27–29 (20–22 in DeRose 2017), apart from 122 n. 30 and 204), doesn't use it here to characterize what the reader could say: "I know the Bulls won but my paper might have reported this incorrectly!". Or is this not so abominable after all? It would be interesting to hear more about this.

DeRose's comparison. Again (146), DeRose briefly mentions the skeptic who wants to raise doubt about knowledge at Level I by pointing to uncertainty at Level II (the less than perfect reliability of methods used).

The main upshot of all this is that it would be good to be clear about the difference between epistemic lotteries and lottery propositions. I wonder whether DeRose would be willing to accept this.

6. Knowing Lottery Propositions

Before we move on to DeRose's contextualist solution of (his version of) the lottery puzzle, it would be good to pause briefly and think a bit—contextualism or no contextualism—about whether the verdict that one cannot know lottery propositions really is intuitive or even true. I will try to stay away as much as possible from using any substantial, controversial philosophical views here.

First of all: why do newspapers print lottery results if all those with losing tickets already know that they have lost? Do they only print the results for the sake of the one or very few winners? That would be weird. However, if what I said above is correct, then knowledge doesn't always terminate inquiry and sometimes (like in cases of outcomes of regular lotteries) there are good reasons to check further and, e.g., read the paper (see Section 4 above).[24]

[24] One of the most popular (and quickest) arguments against the possibility of knowing that one's lottery ticket has lost (just on the basis of probabilistic reasoning) goes like this: if people know in advance that the ticket they're buying is a loser, then they don't have a good reason to buy the ticket. But they do have some good reason to buy the ticket. Hence, they don't know that their ticket is a loser. Taking a closer look at their "good reasons" to buy the ticket reveals that this argument is not very strong. Considerations

It is also useful to use a method of "intermediate cases" (see, e.g., and in a different context, Fischer and Ravizza 1992) here. Consider the following series of cases of lotteries:[25]

(a) a regular 1-Million-ticket lottery (the subject's ticket loses);

(b) like (a) except that all the losing tickets are changed into winning tickets and the winning ticket into a losing ticket (the subject is one of the many winners);

(c) like (b) except that "winning" means staying alive and "losing" means being shot (the subject is not being shot);

(d) like (c) except that the player is playing Russian Roulette with a revolver that has 1 Million chambers but only one bullet in it (the subject doesn't shoot themselves);

(e) like (d) except that the subject is taking medication that has lethal side effects in 1 out of 1 Million cases (the subject experiences no side effects).

The differences between these five cases do not seem relevant at all to the question whether the subject can know in advance the probable outcome of the lottery (if it happens). There is also no bad slippery slope going on here (the differences are irrelevant to questions about knowledge). Hence, if the subject in (a) cannot know that their ticket is a loser, then the subject in (e) cannot know that they will survive taking their medication. I find it very plausible to judge that the

of expected utility speak against buying the ticket. But there are other, good reasons for buying a ticket after all: for instance, the entertaining suspension of disbelief (that one will win) and the joy of day-dreaming.

[25] Nothing depends on the use of the word "lottery" in all these cases.

subject in (e) knows that they will survive taking their medication. Hence, it is also very plausible to judge that the subject in (a) knows that they've lost the lottery.[26]

Arguments from equivalent lotteries can also be made in other ways. Suppose that one person, Mr. Havit, owns[27] 999,999 out of the 1 Million tickets in a lottery; the remaining ticket is owned by Mr. Nogot. Suppose both know this. Finally, suppose that one of Havit's tickets is the unique winner. It is, I find, implausible to a large degree to claim that Havit cannot know in advance that he is winning the lottery. But if Havit can know before the drawing that he's winning the lottery, then Nogot can figure this out, too (how could he not?). If Nogot knows that Havit is winning, then he can come to know, by simple inference,[28] that he himself is losing (both of them know that if Havit is winning, then Nogot is losing). Since it doesn't make a difference to Nogot's epistemic position with respect to the lottery proposition *Nogot's ticket is a loser* whether all the other 999,999 tickets are owned by 1, 2, 11 or 999,999 persons, anyone in the original regular lottery scenario should be able to know that their ticket is losing.

These remarks are meant to support DeRose's similar views and suggest that it is not implausible at all to claim that one can know, merely on the basis of probabilistic evidence, that one's ticket is a loser.

[26] If one doesn't find (e) that convincing, one can add further steps, like this one:

> (f) like (e) except that the subject is crossing the street where there is a 1 in a Million chance to be run over and killed (the subject remains unharmed).

[27] For whatever reason: perhaps Havit is an eccentric billionaire trying to win a bet that he can win this lottery. Eccentricity also prevented him for buying the last remaining ticket, too. For such a case, see Baumann (2016: 107; see also two other cases there on 107–108).

[28] There are no problems with closure here.

7. A Contextualist Solution

DeRose applies the basic idea for his well-known solution to the skeptical puzzle (see his 1995) to the lottery puzzle (section 5.14), even though he doesn't use the analogy between them in the set-up of the problem (see above). He notes that using a sensitivity account of knowledge here would not help because the belief (acquired after reading the paper) that the Bulls won (or that one's ticket is a loser) is sensitive while the belief (before reading the paper and merely based on probabilistic evidence) that one's ticket is a loser is not sensitive. Hence, the first belief can be knowledge while the second cannot. This contradicts the intuitive comparative judgment which DeRose wants to respect: if one can (cannot) know the one, then one also can (cannot) know the other (152–153). And this should turn out true no matter what epistemic standards are in place (154).

Rather, DeRose uses his Rule of Sensitivity, though not as part of an explanation of the concept of knowledge, and applies it in a contextualist way (155). Here is that rule again: "When it's asserted that S knows (or doesn't know) that P, then, if necessary, enlarge the sphere of epistemically relevant worlds so that it at least includes the close worlds in which P is false" (1995: 37; 2017: 28).[29] In ordinary, lower-standard contexts sensitivity is not amongst the standards for knowledge. Hence, in such a context, an assertion of "Keith knows that the Bulls won" can be true; similarly, in such a context, an assertion of "Jackie knows that her ticket is a loser" can be true, too. I am assuming that testimony in the first case and probabilistic reasoning in the second case are sufficient here, according to DeRose. However, when someone asserts (or

[29] One might want to add that thinking without asserting that S knows (or doesn't know) that *p* does the same. The explanation above doesn't go that well with the remark mentioned before about possible worlds in DeRose (2017: 154 n.10).

merely thinks?) that their ticket is a loser, the Rule of Sensitivity is triggered by default (which trigger can be overridden: see sections 5.17 and 5.18) and added to the standards for knowledge. This makes the epistemic context more demanding. Since Jackie's belief that her ticket is a loser (remember: she only has her probabilistic reasoning as support for this) is not sensitive, and since sensitivity is a condition for a true assertion of "Jackie knows that her ticket is a loser" in this more demanding context, the claim "Jackie knows that her ticket is a loser" now turns out to be false. The context-change has changed the truth-conditions for "knows" (155).

Being myself quite sympathetic to contextualism in general as well as to a contextualist response to lottery problems (see Baumann 2016: ch. 4.5–4.6), I won't say much at all about this aspect here (that I favor a different kind of contextualism doesn't matter here). I've already (see Section 3) raised some questions about why sensitivity should play such a crucial role and about how one can respond to certain problems with the sensitivity condition in lottery cases. I won't repeat this. There are further questions to be raised here. First, what about the belief that the Bulls won? DeRose holds that it's sensitive. So, one should assume that it would survive the change to a more demanding context where the Rule of Sensitivity is in force. DeRose, however, wants to hold on to the comparative judgment that if one doesn't know (before reading) that one's ticket is a loser, then one also doesn't know that the Bulls one. So, one open question is: what prevents Keith's belief that the Bulls won from being truly called "knowledge" in a more demanding context? Are there other epistemic standards in play, in addition to the Rule of Sensitivity, which would kick the Bulls-belief off the pedestal of knowledge (or "knowledge") in that more demanding context? Or do more demanding contexts also raise questions at Level II, questions about whether one has won one's epistemic lottery? But why should this be part of a demanding context? An alternative move would be to give up on the comparative judgment; one could, for instance, assume that epistemic standards don't apply equally across the range of propositions but are proposition-specific (see Baumann 2016: 137–138).

How would DeRose's solution work for *Lottery Problem* (see above) about which he doesn't say much? Here is the case of Jackie's Puzzle again:

Jackie's Puzzle

(i) Jackie does not know that (a) she won't win the lottery;

(ii) Jackie does know that (b) she won't be able to afford a much bigger apartment any time soon;

(iii) If Jackie does not know (a), then she does not know (b).[30]

Applying the Rule of Sensitivity here, we would say that in an ordinary context an utterance of (i) is false and one of (ii) is true. Accordingly, in this context it is true to say both that "Jackie does know that she won't win the lottery" and that "Jackie does know that she won't be able to afford a much bigger apartment any time soon." No problem then with (iii). This part of the puzzle would be resolved. In contrast, in a demanding context an utterance of (i) would have to come out true and one of (ii) false. Accordingly, in this context it is true to say both that "Jackie does not know that she won't win the lottery" and that "Jackie does not know that she won't be able to afford a much bigger apartment any time soon." If we accept this second part of the solution, too, then there is, again, no problem with (iii) and the puzzle seems to be completely resolved. However, one question, at least, remains: What (in a demanding context) prevents a claim that "Jackie does know that she won't be able to afford a much bigger apartment any time

[30] We can even strengthen "does not know" to "cannot know".

soon" from being true (see a brief statement without further argument on this on 162)? Apparently, not the Rule of Sensitivity. Again, one can have the impression that one crucial element in DeRose's solution to the (his) lottery problem is still missing.

One last observation on all this. I mentioned the possible view that epistemic standards for knowledge are proposition specific. This would raise serious questions about comparative judgments of the type DeRose is referring to. Apart from this complication, there is the further question whether someone couldn't have extraordinarily good evidence for some ordinary proposition *o* (I can see that my car is parked on F37, I hold the picture I've taken of the slot in my hands, etc.) and not so great evidence for some lottery proposition which is entailed by the ordinary proposition *l* (that I'm not suffering from a rare "car on F-37"-hallucination while my car is being stolen and taken away from F-37). Suppose that we do have proposition-independent epistemic standards after all. Then it seems that nothing stands in the way of imagining some not too demanding and not too lax context in which it is true to say that the subject knows that *o* but does not know that *l*. This is a problem if one wants to hold on to some principle of epistemic closure (like (Closure) above) and also wants to assume that our subject knows about the entailment of *l* by *o* or can or does infer *l* from *o*.[31]

8. A Final Problem for Knowers of Lottery Propositions: Absurdity?

DeRose also addresses what is perhaps the most serious objection against anyone (contextualist or not) who wants to allow for knowledge of lottery propositions. If I can know of each losing ticket that it's losing, then I can also come to know, by inference, that all these (losing) tickets

---

[31] For a modified principle of closure that could, perhaps, deal with this problem and also be of use to the contextualist, see Baumann (2016: ch. 4.4 and also ch. 4.1).

are losing. This is already pretty absurd. By further simple inference, I can then also come to know which ticket is the winner—which adds further to the absurdity (DeRose 2017: 161; for this objection see, e.g., Nelkin 2000: 373–374; Hawthorne 2002: 246–247, and 2004: 6–7, 15–16, 19–20, 47–49, 94–98, 145–146, as well as Goldman 2008: 466).

I do find DeRose's reply to this objection very convincing (and will therefore not go much into the details of it): the objection assumes that one can just move from "each" to "all", or, in other words, that just on the basis of knowing each of the many conjuncts (*Ticket 1 is a loser*, *Ticket 2 is a loser*, … *Ticket n-1 is a loser*) one can infer and thus come to know the conjunction (*Ticket 1 is a loser*, and *Ticket 2 is a loser*, … and *Ticket n-1 is a loser*); from the latter one can then, of course, infer that ticket n is the winner and all others losers. The step from knowledge of each of the many conjuncts to the conjunction presupposes a principle of multi-premise closure, a simple version of which (which should suffice here[32]) could be formulated in the following way:

> (Multi-Premise Closure) If S knows that *p1*, knows that *p2*, …, knows that *pn*, and if S also knows that *p1*, *p2*, …, *pn* entail their conjunction, then S knows the conjunction *p1, and p2, ... and pn*.[33]

DeRose argues (see his sections 5.20–5.22), convincingly, that (Multi-Premise Closure) is not plausible and way too strong for an acceptable closure principle. One should reject it and restrict

[32] See n. 6 on the sister principle (Closure).

[33] See also Hawthorne (2004: *passim*). I am putting questions concerning the individuation of premises aside here.

one's closure principles accordingly to something like (Closure). This alone blocks the absurdity objection (see also Baumann 2016: 195–207). I would like to add some further points to the defense.

Consider, for instance, this case. Daniel is very good at identifying dachshunds. He has seen 5000 animals 4999 of which he has correctly identified as dachshunds, whereas the one remaining animal (an extremely rare animal called "dachsfox") he misidentified as a dachshund. It seems very implausible to deny that of each individual dachshund Daniel knew that it is a dachshund. In each individual case, he was pretty certain though not perfectly certain and his justification or warrant was pretty good but not maximally good. In analogy to the fact that a conjunction of many conjuncts with high but not maximal probability will itself have a much lower probability, we can say that Daniel's epistemic position with respect to the conjunction of the 4999 true conjuncts is much weaker than his epistemic position with respect to each conjunct, so much weaker that it is not strong enough for knowledge of the conjunction (for doubts about multi-premise closure, see also Olin 2005 and Lasonen-Aarnio 2008).[34] Multi-premise closure is independently implausible.

A comparison of the lottery case with the preface paradox also throws some light on all this. We can easily imagine a reasonable author who has very good reasons to claim (in a preface) that not everything they are saying in the book is true. Reasonable people and authors know that they are fallible. At the same time, such a reasonable author might also have very good reasons for each individual claim that they are making. If this committed them to accepting

[34] Or consider an "inverse fake barn case" (Goldman 1992: 86): someone travels through a part of the country with a lot of barns and exactly one fake barn. Don't they know they're facing a real barn when they do?

the conjunction of all claims made in the book, then our reasonable author would be inconsistent (see Makinson 1965). We already have good reasons to reject a principle of multi-premise closure. Now, in addition we can also easily imagine that our author knows each of the true propositions claimed by him to be true. He does, of course, not know the false propositions. Suppose, finally, that the author's book is just a list of 1 Million claims, each about a different one of the 1 Million tickets in a regular lottery. If we were happy to attribute knowledge of the many true claims in the book before this last modification of our case, then we should also be happy to attribute knowledge here. And similarly for the ordinary lottery case. If the move from "each" to "all" is not convincing in the preface-paradox case, then it should also not be convincing in the regular lottery case. I don't think many people would accept an analogue of the absurdity objection in the preface-paradox case.

I want to add one more response to the absurdity objection: the inference from each of the conjuncts to the conjunction would not be rational and could thus not lead to knowledge (given that knowledge excludes relevant irrationality or lack of rationality). Suppose I believe of each of the n tickets in a lottery that it's a loser. In order to make the inference to the conjunction, I would have to pick n-1 tickets out of the n tickets (assuming we have exactly one winner). Since I don't know which ticket is the winner (I believe of each ticket that it is a loser and thus believe of the winner falsely that it's a loser), I can only make an arbitrary selection of some n-1 tickets (out of the n possible combinations of n-1 tickets). If I pick, by sheer accident, just all of the losers, I would be very lucky and my choice would not be based on reason. Given this lack of rationality, I cannot acquire knowledge of the conjunction.[35]

9. Conclusion

I have found that one should agree on many points with DeRose's impressive treatment of the lottery puzzle. On some issues, I have offered friendly suggestions. On others, I have raised questions. On still others, I have offered criticism (the lack of engagement with *Lottery Problem* or the missing distinction between "levels" of lotteries, that is, between lottery propositions and epistemic lotteries). For dramatic effect some might sometimes prefer more deep disagreement and more confrontational discussion about some "big" questions. However, in the case of DeRose's complex and multi-faceted *The Appearance of Ignorance*, going a bit more into the details is really much more interesting and useful, especially if one agrees with much of it.

References

Adler, Jonathan E. 1986. "Knowing, Betting and Cohering," *Philosophical Topics* 14: 243–257.

Baumann, Peter. 2011. "Epistemic Closure." In Sven Bernecker & Duncan Pritchard (eds.), *The Routledge Companion to Epistemology*, 597–608. London: Routledge.

Baumann, Peter. 2016. *Epistemic Contextualism. A Defense*. Oxford: Oxford University Press.

Baumann, Peter. MS. "Remarks on Stakes."

Bondy, Patrick. 2013. "How to Understand and Solve the Lottery Paradox," *Logos & Episteme*

[35] As a side remark, one can add that in a standard case I know of each individual ticket that it's a loser but I don't know that I know that it's a loser (otherwise I could, indirectly, identify the winner). KK-principles according to which one knows that one knows that *p* if one knows that *p* are very problematic.

4: 283–292.

Brown, Jessica. 2008. "Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning," *Noûs* 42: 167–189.

Brueckner, Anthony. 2012. "Roush on Knowledge: Tracking Redux?" In Kelly Becker & Tim Black (eds.), *The Sensitivity Principle in Epistemology*, 229–241. Cambridge: Cambridge University Press.

Cohen, Stewart. 1988. "How to Be a Fallibilist," *Philosophical Perspectives* 2: 91–123.

Collins, John. 2006. "Lotteries and the Close Shave Principle." In Stephen Hetherington (ed.), *Aspects of Knowing. Epistemological Essays*, 83–96. Amsterdam: Elsevier.

DeRose, Keith. 1995. "Solving the Skeptical Problem," *The Philosophical Review* 104: 1–52.

DeRose, Keith. 1996. "Knowledge, Assertion and Lotteries," *Australasian Journal of Philosophy* 74: 568–580.

DeRose, Keith. 2017. *The Appearance of Ignorance. Knowledge, Skepticism and Context, Volume 2*. Oxford: Oxford University Press.

Douven, Igor. 2007. "A Pragmatic Dissolution of Harman's Paradox," *Philosophy and Phenomenlogical Research* 74: 326–345.

Dudman, V.H. 1992. "Probability and Assertion," *Analysis* 52: 204–211.

Ebert, Philip A., Smith, Martin & Durbach, Ian. 2018. "Lottery Judgments: A Philosophical and Experimental Study," *Philosophical Psychology* 31: 110–138.

Fischer, John M. & Ravizza, Mark 1992. "Thomson and the Trolley," *Journal of Social Philosophy* 23: 64–87.

Gettier, Edmund L. 1963. "Is Justified True Belief Knowledge?," *Analysis* 23: 121–123.

Goldman, Alan. 2008. "Knowledge, Explanation, and Lotteries," *Noûs* 42: 466–481.

Goldman, Alvin I. 1992. "Discrimination and Perceptual Knowledge." In his *Liaisons.*

Philosophy Meets the Cognitive and Social Sciences, 85–103. Cambridge, MA & London: MIT Press.

Greco, John. 2003. "Knowledge as Credit for True Belief." In Michael DePaul & Linda Zagzebski (eds.), *Intellectual Virtue: Perspectives from Ethics and Epistemology*, 111–134. Oxford: Oxford University Press.

Greco, John. 2004. "A Different Sort of Contextualism," *Erkenntnis* 61: 383–400.

Harman, Gilbert. 1968. "Knowledge, Inference, and Explanation," *American Philosophical Quarterly* 5: 164–173.

Harman, Gilbert. 1973. *Thought*. Princeton: Princeton University Press.

Harman, Gilbert & Sherman, Brett. 2004. "Knowledge, Assumptions, Lotteries," *Philosophical Issues* 14: 492–500.

Hawthorne, John. 2002. "Lewis, the Lottery, and the Preface," *Analysis* 62: 242–251.

Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Clarendon.

Kyburg, Henry E. Jr. 1961. *Probability and the Logic of Rational Belief*. Middletown: Wesleyan University Press.

Lasonen-Aarnio, Maria. 2008. "Single Premise Deduction and Risk," *Philosophical Studies* 141: 157–173.

Lewis, David. 1996. "Elusive Knowledge," *Australasian Journal of Philosophy* 74: 549–567.

Makinson, D.C. 1965. "The Paradox of the Preface," *Analysis* 25: 205–207.

McKinnon, Rachel. 2013. "Lotteries, Knowledge, and Irrelevant Alternatives," *Dialogue* 52: 523–549.

Mills, Eugene. 2012. "Lotteries, Quasi-Lotteries, and Scepticism," *Australasian Journal of Philosophy* 90: 335–352.

Nagel, Jennifer. 2011. "The Psychological Basis of the Harman-Vogel Paradox," *Philosophers'*

Imprint 11.5: 1–28.

Nelkin, Dana K. 2000. "The Lottery Paradox, Knowledge, and Rationality," *The Philosophical Review* 109: 373–409.

Olen, Jeffrey. 1977. "Knowledge, Probability, and Nomic Connections," *Southern Journal of Philosophy* 15: 521–526.

Olin, Doris. 2003. *Paradox*, Chesham Bucks: Acumen.

Olin, Doris. 2005. "A Case against Closure," *Veritas* 50: 235–247.

Roush, Sherrilyn. 2005. *Tracking Truth. Knowledge, Evidence, and Science*. Oxford: Clarendon.

Roush, Sherrilyn. 2012. "Sensitivity and Closure." In Kelly Becker & Tim Black (eds.), *The Sensitivity Principle in Epistemology*, 242–268. Cambridge: Cambridge University Press.

Ryan, Sharon 1996. "The Epistemic Virtue of Consistency," *Synthese* 109: 121–140.

Stemmer, Nathan. 1982. "A Solution to the Lottery Paradox," *Synthese* 51: 339–353.

Timmerman, Travis. 2013. "The Persistent Problem of the Lottery Paradox and its Unwelcome Consequences for Contextualism," *Logos & Episteme* 4: 85–100.

Turri, John & Friedman, Ori. 2014. "Winners and Losers in the Folk Epistemology of Lotteries." In James R. Beebe (ed.), *Advances in Experimental Epistemology*, 45–69. London: Bloomsbury.

Vogel, Jonathan 1990. "Are there Counterexamples to the Closure Principle?" In Michael D. Roth & Glenn Ross (eds.), *Doubting. Contemporary Perspectives on Skepticism*: 13–27. Dordrecht: Kluwer.

Weintraub, Ruth 2001. "The Lottery: A Paradox Regained and Resolved," *Synthese* 129: 439–449.