

# THE PERSISTENT PROBLEM OF THE LOTTERY PARADOX: AND ITS UNWELCOME CONSEQUENCES FOR CONTEXTUALISM

Travis TIMMERMAN

**ABSTRACT:** This paper attempts to show that contextualism cannot adequately handle all versions of “The Lottery Paradox.” Although the application of contextualist rules is meant to vindicate the intuitive distinction between cases of knowledge and non-knowledge, it fails to do so when applied to certain versions of “The Lottery Paradox.” In making my argument, I first briefly explain why this issue should be of central importance for contextualism. I then review Lewis’ contextualism before offering my argument that the lottery paradox persists on all contextualist accounts. Although I argue that the contextualist does not fare well, hope nevertheless remains. For, on Lewis’ behalf, I offer what I take to be the best solution for the contextualist and argue that once this solution is adopted, contextualism will be in a better position to handle the lottery paradox than any other substantive epistemological theory.

**KEYWORDS:** lottery paradox, contextualism, epistemology

## I. The Lottery Paradox

There are a few epistemological puzzles that revolve around what is referred to as “The Lottery Paradox.” A standard formulation of one of these paradoxes is as follows. Suppose Poor Bill and Skeptical Susan are talking and Susan invites Bill to come with her on an African Safari next year. Bill, a wage slave, politely declines, saying that he will not have enough money to go. Now, suppose that Bill plays the lottery each week and if he were to win the lottery, he would have enough money to go on an African Safari. If Bill knows that he will not have enough money to go on the safari and Bill recognizes that this entails that he will lose the lottery, then by the closure principle,<sup>1</sup> Bill knows that he will lose the lottery; hence the paradox. It seems that while Bill knows he will not have enough money to go on a

---

<sup>1</sup> The closure principle can be formulated as “If person *S* knows *p*, and *p* entails *q*, then *S* knows *q*.” This definition is from Steven Luper, “The Epistemic Closure Principle,” *The Stanford Encyclopedia of Philosophy (Fall 2012 Edition)*, Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2012/entries/closure-epistemic/>>.

safari next year, he does *not* know that he will never win the lottery. How can this be?

David Lewis, and other contextualists,<sup>2</sup> attempt to solve the paradox by arguing that the truth-value of these knowledge ascriptions are sensitive to certain facts about the context in which they are uttered. The facts about context that are considered relevant differ between contextualists, so for simplicity's sake, I will focus on those which Lewis gives in his seminal work, "Elusive Knowledge."<sup>3</sup> If Lewis' contextualism can be made to work by its own lights, then it needs to account for our intuition that knowledge claims like "I know I will lose the lottery" are false, while maintaining that we can still rightly claim to know propositions such as "I know I will not have enough money to go on an African Safari next year."<sup>4</sup> Unfortunately, Lewis' account fails to do just that, or so I will argue. At the same time, Lewis' contextualist solution is ingenious and unique. One of its many virtues is that it allows for us to know that we will lose the lottery when the *Rule of Resemblance* is not salient, which, as we will see, is the intuitively right result.<sup>5</sup> Yet, once we consider how the *Rule of Resemblance* applies in analogue lottery cases, Lewis becomes stuck between a rock and a hard place. Specifically, he will either have to deny we have knowledge in cases where it intuitively seems like we have knowledge, or grant that we *can* know we will lose the lottery in contexts in which we seem to lack knowledge about whether we will lose the lottery. Hence, if my argument works, there will be two horns that Lewis will have to choose from. Either we do not know that we will lose the

---

<sup>2</sup> For similar contextualist accounts that could be used, please see Stewart Cohen, "How to be a Fallibilist," *Philosophical Perspectives* 91 (1988): 581-605, Stewart Cohen, "Skepticism, Relevance, and Relativity," in *Dretske and His Critics*, ed. Brian McLaughlin (Massachusetts: Blackwell Press, 1991), 17-37, and especially Stewart Cohen, "Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery," *Australasian Journal of Philosophy* 76 (1998): 289-306. See also Peter Unger, "The Cone Model of Knowledge," *Philosophical Topics* 14 (1986): 125-178, Peter Unger, *Philosophical Relativity* (Minneapolis: University of Minnesota Press, 1984), and Keith DeRose, "Solving the Sceptical Problem," *Philosophical Review* 104 (1995): 1-5. Some of Cohen's work uses his fallibilism to solve the lottery paradox in a way similar to, but less complex than, Lewis' solution. I use Lewis' account in this paper because I believe it's the strongest form of contextualism. I focus only on his contextualist rules for both simplicity's sake and length issues. Suffice it to say that alternative accounts are similar enough to Lewis' that they do not seem to be able to avoid the objections I raise in this paper. Lewis' account seems to get the *least* wrong in lottery cases.

<sup>3</sup> David Lewis, "Elusive Knowledge," *Australasian Journal of Philosophy* 74 (1996): 549-567.

<sup>4</sup> This example is drawn from John Hawthorne, *Knowledge and Lotteries* (New York: Oxford University Press, 2004), 160-162.

<sup>5</sup> Lewis, "Elusive Knowledge," 282.

lottery, but at the expense of something close to universal skepticism, or we can actually know that we will lose the lottery, even when it's salient that we are holding a ticket in a fair lottery.

My paper takes the following form. First, I briefly explain why this issue should be of central importance for contextualism, indeed any epistemological theory, instead of something merely tangential. I then review Lewis' contextualism before offering my argument that the lottery paradox persists on all contextualist accounts, including Lewis' contextualism. Although I argue that the contextualist does not fare well, hope nevertheless remains. For, on Lewis' behalf, I offer what I take to be the best solution for the contextualist and argue that contextualism is in a better position to handle the lottery paradox than any other substantive epistemological theory. I end the paper with a brief digression, examining how contextualists can handle another formulation of the lottery paradox that concerns *the sufficiency thesis* and the *conjunction principle*.

## II. The Importance of Addressing the Lottery Paradox

Lottery paradoxes may seem like a relatively minor issue in epistemology. Whether an epistemological theory can account for our intuitions in lottery cases seems less crucial than whether it is consistent, can avoid skepticism, captures *most* of our intuitions about which knowledge ascriptions are accurate and handles relevantly similar issues. If lottery cases were isolated components of all epistemological theories, I would agree. However, lottery cases are of central importance to any substantive epistemological theory precisely because the way lottery cases are dealt with has important implications for each of the aforementioned aspects of any epistemological theory. In other words, lottery cases which are not properly accounted for run the risk of being generalized within a theory.<sup>6</sup> Generalizing the rules that apply to lottery cases usually exposes inconsistency with the theory in question. Revising the theory in light of the inconsistency can often result in undermining many knowledge claims to which we feel entitled. This is what I take the issue to be with contextualism generally, and Lewis' account, specifically. As such, examining how an epistemological theory handles lottery paradoxes seems to be of crucial importance.

---

<sup>6</sup> For more on this issue, see Igor Douven and Timothy Williamson, "Generalizing the Lottery Paradox," *The British Journal for the Philosophy of Science* 57 (2006): 755-779.

### III. (Lewis') Contextualism

In an attempt to find a middle ground between two (supposedly) undesirable epistemological theories (i.e. skepticism and fallibilism), Lewis opts for a contextualist framework, which *essentially* consists of five rules combined with a definition of knowledge. These rules are the *Rule of Actuality*, the *Rule of Belief*, the *Rule of Resemblance*, the *Rule of Reliability* and the *Rule of Attention*. It might be useful to start by giving a definition of knowledge and building upon that. Lewis can assert that a subject "S knows that P iff, for every possibility W in which not-P, S knows that not-W" and then add a detailed contextualist framework.<sup>7</sup> The scope of possibilities in this definition is restricted to those possibilities which may not properly be ignored, and the possibilities that may not be properly ignored are determined by context. How might one determine the relevant role of context? Lewis' five rules are supposed to provide the way to distinguish between those possibilities which may be properly ignored and those which may not. The consequence is that one is able to maintain her ordinary everyday knowledge (e.g. I have hands) most of the time. She only fails to know these claims once the context shifts, preventing one (or more) of the five rules from being met.

In what follows, I will review Lewis' five contextualist rules while explaining how his contextualism is supposed to handle a formulation of the lottery paradox. Lewis' contextualist solution is original and *prima facie* plausible. His account allows for us to know that we will lose the lottery when the Rule of Resemblance is not salient.<sup>8</sup> But in any context where this rule becomes salient, we will lose knowledge that we won't win the lottery. Applying Lewis' contextualist rules to lottery cases will yield the right result in most, but not all, cases. Before I review what I take the problematic cases to be, I will offer an exposition of Lewis' contextualist rules. The first rule is the *Rule of Actuality* and is simply the stipulation that the "possibility that actually obtains is never properly ignored."<sup>9</sup> This should be fairly straightforward and accounts for the truth condition of knowledge. I cannot know that I will lose the lottery if I have the winning ticket.<sup>10</sup> The second rule is the *Rule of Belief* and is also fairly

---

<sup>7</sup> Lewis, "Elusive Knowledge," 273.

<sup>8</sup> Lewis, "Elusive Knowledge," 282.

<sup>9</sup> Lewis, "Elusive Knowledge," 274.

<sup>10</sup> It might be worth noting that this rule is an externalist one. That is, we will almost never (and never with skeptical hypotheses) be able to determine with absolute certainty whether this rule is met. As such, we might not have meta-knowledge in many cases, which may be an unwelcome conclusion for some.

straightforward. It is the claim that a “possibility that the subject believes to obtain is not properly ignored” and this is true “whether or not he is right to so believe.”<sup>11</sup> Thus, I cannot know that I have hands if I actually believe that I am a BIV and consequently believe that I do not have hands. This just accounts for the belief condition of knowledge.

The third rule is the *Rule of Resemblance*, which is a bit more complicated. Here is the most straightforward and concise manner in which it can be stated: If two possibilities saliently resemble one another and “if one of them may not be properly ignored, neither may the other.”<sup>12</sup> This rule is tricky, and the trickiness occurs as a result of the qualifier ‘salient,’ as well as the ambiguity of how the term ‘resemble’ is being used.<sup>13</sup> Although Lewis never provides an explicit account of how these terms are being used, we can avoid any problems of ambiguity by considering clear examples on both ends. I will do this shortly. A final point about the rule is worth noting. Lewis acknowledges that there is an *ad hoc* element to its application.<sup>14</sup> It is not applied to the resemblance that any skeptical possibility resembles actuality with respect to the subject’s evidence.<sup>15</sup> For if it were applied in that way, then (near) universal skepticism would be the result. Lewis appeals to the *Rule of Resemblance* to take care of the lottery problem.<sup>16</sup> We will therefore return to it shortly.

The final three rules are the *Rule of Reliability*, the *Rule of Conservatism* and the *Rule of Attention*. The *Rule of Reliability* is exactly what it sounds like. It requires that knowledge be obtained by a reliable process (e.g. vision). If you acquire a true belief that meets the other rules by some unreliable process (e.g. palm reading), then you lack knowledge.<sup>17</sup> The *Rule of Conservatism* allows (very defeasibly) that we may properly ignore what those around us ignore.<sup>18</sup> This will also play a role in the lottery paradox discussion. The *Rule of Attention* is almost

---

<sup>11</sup> Lewis, “Elusive Knowledge,” 275.

<sup>12</sup> Lewis, “Elusive Knowledge,” 275.

<sup>13</sup> There are also issues Lewis notes, such as cases where “one possibility saliently resembles two or more others,” where one resembles the second in one respect, but resembles the third possibility in another respect.

<sup>14</sup> Lewis, “Elusive Knowledge,” 276.

<sup>15</sup> Lewis, “Elusive Knowledge,” 276.

<sup>16</sup> It’s worth noting that the *Rule of Resemblance* applied to the lottery paradox is similar to Cohen’s ‘salience’ rule. See Cohen, “How to be a Fallibilist,” 121 and Hawthorne, *Knowledge and Lotteries*, 159.

<sup>17</sup> Lewis, “Elusive Knowledge,” 277. This rule is, of course, defeasible. The *Rule of Actuality* alone or conjoined with the *Rule of Resemblance* can easily undermine the *Rule of Reliability*.

<sup>18</sup> Lewis, “Elusive Knowledge,” 277. Again, this is defeasible and could be undermined by any of the other rules.

tautological. It simply requires that whatever is properly ignored is *as a matter of fact* being ignored. It's not enough for it to be the case that it *could* be properly ignored by the individual. It also has to actually be the case that it *is currently being ignored by the individual*.<sup>19</sup>

#### IV. How Contextualists Try to Handle the Lottery Paradox

At this point, we can consider how Lewis' account handles the lottery paradox. The *Rule of Resemblance* is supposed to rule out any case of knowing I will lose the lottery when I am thinking about the lottery.<sup>20</sup> Lewis argues that for "every ticket, there is the possibility that it will win," which means that these "possibilities are saliently similar to one another: so either every one of them may be ignored, or else none may."<sup>21</sup> But one of them will be the winning ticket, so by the *Rule of Actuality*, it may not properly be ignored. Since one may not be properly ignored, and since they all saliently resemble one another, none may be properly ignored.<sup>22</sup> Now, while we cannot properly ignore the possibility of having the winning ticket when we are thinking about playing the lottery, we can (according to Lewis) properly ignore the possibility of having the winning lottery ticket when we are properly ignoring the fact that we are (or could in the future) play the lottery. Consider Lewis' case of Poor Bill again. Poor Bill is a wage slave who spends all of his spare cash gambling, including playing the lottery. We might say that we know "Poor Bill will never be rich." But if the possibility that Bill's ticket wins saliently resembles the actual winning ticket, then we cannot properly ignore the possibility that Bill will win the lottery and therefore would not know that he will never be rich. But Lewis' account contains a loophole due to his qualifier 'salient.' Lewis writes that when talking about the cases in which one is considering the fact that she is playing the lottery ...

I saw to it that the resemblance between the many possibilities associated with the many tickets was sufficiently salient. But this time, when we were busy pitying poor Bill for his habits and not for his luck, the resemblance of the many possibilities was not so salient. At that point, the possibility of Bill's winning was properly ignored; so then it was true to say that we knew he would never be rich

---

<sup>19</sup> Lewis, "Elusive Knowledge," 277-278.

<sup>20</sup> Specifically, Lewis states "It is the *Rule of Resemblance* that explains why you do not know that you will lose the lottery, no matter what the odds are against you and no matter how sure you should therefore be that you will lose." (277).

<sup>21</sup> Lewis, "Elusive Knowledge," 276.

<sup>22</sup> Lewis, "Elusive Knowledge," 276.

... And at that point, it was also true that we knew he would lose – but that was only true so long as it remained unsaid! (And maybe unthought as well).<sup>23</sup>

If this works, Lewis provides us with a nifty solution to this version of the lottery paradox, but I have reservations about whether it can work given (i) Lewis' stipulation that there is necessarily a winner in his lottery cases and (ii) his 'salience' qualifier. If my argument against Lewis is sound, then Lewis' account (as described) will lead to a much wider skepticism than he (or any contextualist) would want to embrace. The first thing to note is that the 'salience' qualifier can be applied to lottery cases too, as which facts and possibilities are salient in any individual's mind is merely a contingent matter. It's a psychological fact that when *most* people play the lottery, the possibility that they will win is salient, but it's not a logical truth that when one plays the lottery the possibility of winning is salient to the one playing. There are possible worlds where people play fair lotteries and the fact that there is a chance they will win is never salient to any player. John Hawthorne even notes that some people in the actual world play the lottery without the possibility of winning ever becoming salient to them.<sup>24</sup>

## V. Some Problems with Lewis' Contextualist Account

I will now consider an issue with (i), consider the best response on Lewis' behalf and then argue that this response does not avoid the lottery paradox without resulting in a kind of skepticism. In the process of doing this, I will discuss (ii) as well. Suppose, for the sake of argument, that Lewis' account works for lottery cases where there is a guaranteed winner. The following question naturally arises. What about cases where no one is guaranteed to win? In the actual world, there are lotteries with no guaranteed winner and most people have the intuition that one *cannot* know she will lose *any* fair lottery (regardless of whether there is a guaranteed winner).

---

<sup>23</sup> Lewis, "Elusive Knowledge," 282.

<sup>24</sup> Hawthorne, *Knowledge and Lotteries*, 84. At least, he seems to say something along these lines. In his discussion of fallibilism's handling of the lottery paradox, he responds to Cohen's salience criteria. He asks "Why should it be inevitable that I take the falsity of the lottery proposition seriously once the question is raised?" The problem is that despite this fact that it may be plausible to claim that the minute chance of winning the lottery may be salient in most cases, it is not necessarily so for every case. Indeed, Hawthorne notes that sometimes people do "flat-out assert that they will not win the lottery." Maybe one cannot assert that he will lose the lottery without the possibility that he will win being salient to him. If so, we can just revise the example to avoid this problem. My own example is this section should get around this objection.

There are two possible answers Lewis could give in lottery cases with no guaranteed winner. On the one hand, he could just accept that we can know we will lose the lottery in such cases. But he probably does not want to take that route for a few reasons. First, it fails to capture the intuition that people have about lotteries, which is an issue contextualism is supposed to solve. If Lewis' contextualism can only show that it gets some subset of lottery cases right and others wrong, it fails to achieve this goal. Second, the distinction between the lottery cases where there is a guaranteed winner and one where there is not seems *ad hoc* and therefore unjustified. It's *ad hoc* because the probability that one holds a winning ticket could be identical in both lottery cases and any combination of Lewis' contextualist rules could be met or violated in either lottery case.

As an illustration of this point, note that the probability that someone is a winner in a fair lottery (with a guaranteed winner) is determined by the number of tickets held over the number of tickets there are. If someone has one ticket, then her odds of winning are  $1/n$ , where  $n$  is the number of tickets being held in the lottery. In a fair lottery where there is no guaranteed winner, the probability that an individual will win is determined by the number of tickets she has over the number of total possible combinations of lottery numbers. So, if an individual has one ticket, then her chances of winning would be  $1/c$ , where  $c$  is the total number of possible winning number combinations in the lottery. The odds of any fair lottery *with* a guaranteed winner ( $1/n$ ) can be made identical to the odds of any fair lottery *without* a guaranteed winner ( $1/c$ ). We only need to make the number of tickets ( $n$ ) in the lottery (with a guaranteed winner) equal to the number of total winning number combinations ( $c$ ) in the lottery (without a guaranteed winner). Then, we have a case where the probability of winning or losing is the same between the two lottery options, each lottery is fair (i.e. there is no trickery going on in selecting the winner or winning numbers), the possibility of winning is not salient and any combination of Lewis' rules could either be met or violated. Since both types of fair lotteries bear all of the relevant similarities, it is seemingly absurd for Lewis to ascribe knowledge in the case where there is no guaranteed winner and deny it in the case where there is one.<sup>25</sup>

---

<sup>25</sup> Some people might still have the intuition that there is some relevant difference between the two lottery cases. They might argue that it's the fact that someone is guaranteed to win that is relevant. There are two things to say in response. First, it's not clear why the fact that someone will win is relevant. The odds of winning for each ticket holder is the same in both cases. The common intuition seems to be that a person does not know she will lose the lottery *because* they might be the winner regardless of whether someone else is guaranteed to win if they lose. Second, even if this is a relevant difference, Lewis' contextualism would still fail to capture our intuition in lottery cases where there is no guaranteed winner. I could then reframe my



Fortunately, one need not make such *ad hoc* distinctions. Since both lottery types are relevantly similar, contextualists should want to treat them as such. They can do this by arguing that we do *not* know that we will lose the lottery in either case. But then a new, equally troubling, question arises. Can Lewis' contextualist rules be applied to exclude knowledge of lottery cases where there is no guaranteed winner? Unfortunately not, for it looks like revising the rules to account for this type of lottery case will inevitably result in the issue of salience coming back to haunt Lewis' account. How might Lewis respond? He could presumably broaden the *Rule of Resemblance* by appealing to counter-factuals. That is, even in cases of fair lotteries with no guaranteed winner, Lewis might argue that the *Rule of Resemblance* still holds for each individual ticket because in some nearby world, there is an individual who had the winning ticket (i.e. the ticket with the correct combination of numbers) and this would clearly be true for every fair lottery. Suppose Lewis does say that. We could then ask him about the following case.

**The Yankees Game** – I pick up the New York Times and read that the Yankees won their last game 6 to 5. I seem to acquire knowledge that the Yankees won 6 to 5, despite the small possibility that there is a typo in the paper.<sup>26</sup>

Yet, if this is true, why could I not similarly acquire knowledge that I will lose the lottery when the odds that I will lose are the same as the odds that there is no typo in the paper?<sup>27</sup> On this interpretation, Lewis cannot appeal to the *Rule of Actuality* to distinguish these two cases, as the lottery case without a guaranteed winner might not have anyone that actually won. Furthermore, it seems like we can appeal to nearby possible worlds in the exact same way in The Yankees Game case as the lottery case. So, broadening the *Rule of Resemblance* to apply to the lottery case (without a guaranteed winner) undermines knowledge in cases like The Yankees Game.

The no-guaranteed-winner lottery case generalizes, and its consequences are far-reaching. There is nothing distinctive about typos in a paper or baseball games. All that was necessary to get the Yankees Game thought experiment off the ground was that (1) It relied on some action (e.g. reading a paper) via which

---

Yankees Game case to mirror this type of lottery case (i.e. It would be a case where the reader is aware that there is some paper out there with a typo in it).

<sup>26</sup> Assume for the sake of argument that I have not yet checked any other sources.

<sup>27</sup> The odds that there is a typo could be quantified in a few different ways. One way would be to divide the number of sentences with typos made by the NY Times in its history by the total number of sentences in the paper's history. Or, if there is one specific paper with a typo in it, we could divide 1 (that paper) over the total number of papers issued that day.

one intuitively gains knowledge and (2) there is some non-zero chance that the person is mistaken about the proposition she comes to believe. And here is where the real problem lies. Almost every knowledge claim we make meets conditions (1) and (2). So, for every knowledge claim we make that meets these two conditions, there is some world with a (no-guaranteed-winner) lottery, where the odds that one is mistaken in her knowledge claim  $p$  is identical to the odds that a person will win the fair lottery in this possible world.<sup>28</sup> Yet, we intuitively want to say that we can truly make these knowledge claims, but deny that the individuals in the possible world know that they will lose the lottery. The distance between the actual world in the lottery case and the nearest possible world where someone wins should be equal to the distance between the actual world where any knowledge claim (meeting conditions (1) and (2)) is made and the nearest possible world where that proposition is false.<sup>29</sup>

Lewis is not without recourse yet. He might rightly claim that winning the lottery when we buy a lottery ticket is salient to us, while entertaining the possibility of a typo in a paper as prestigious as the New York Times is not normally salient to us. For that reason, he might argue that his 'salience' qualifier in the *Rule of Resemblance* would rule out knowledge in the lottery cases, but not in cases like The Yankees Game (or more broadly, any knowledge claims that meet conditions (1) and (2)). Nevertheless, the problem looms. When this rule is applied in the actual world, it is probably an accurate description of *most* people's psychology (when they are buying lottery tickets and reading newspapers), but it is not a necessary truth. In other words, the fact that winning the lottery is salient to most individuals, whereas there being a typo in the New York Times is not salient to most individuals, is merely a contingent matter of fact and so would not intuitively demarcate cases of knowledge from cases of non-knowledge in nearby possible worlds. Moreover, it probably does not get the intuitively right results in our world in every instance. To illustrate this possibility, consider the following case.

---

<sup>28</sup> I'm stipulating that in this possible world, the possibility of winning is not salient in their mind. This should allow all five of Lewis' knowledge-vetoing rules to be avoided.

<sup>29</sup> This is because the odds that the proposition in question is false is always identical between the lottery case and the non-lottery proposition case. It is built into the thought experiment that the world with a typo is a nearby, relevantly similar one. Otherwise, we run into problems with determining what the odds are that there is a typo in a paper in some far away possible world. We can assume that the world is very much like ours and typos occur for the same type of reasons, so the probability that there is a typo is approximately the same between the actual world and any nearby, relevantly similar possible world.

**The Disinterested Lottery Player** – Susie believes that she knows she will lose the lottery each time she plays. However, she also believes that promises confer moral obligations. Susie promised her mother on her deathbed that she would play the lottery each week. Susie’s mother was irrational and played the lottery as often as she could afford a ticket. Sure that her family was bound to win someday, her dying wish was for Susie to play the lottery on her behalf once she was dead. To meet (what she believes is) her moral obligation, Susie does play the lottery weekly for her mother. But, since Susie believes she knows she will lose, each week Susie buys a lottery ticket and then immediately throws it out.

In this case, for Susie, the chance that she will win the lottery is no more salient to her than the chance that a typo in the New York Times will cause her to falsely believe that the Yankees won their last game. That is to say, neither possibility is salient for Susie. Lewis’ contextualism<sup>30</sup> seems to entail that Susie knows she will lose the lottery each week, as none of his rules are violated. Perhaps Lewis would respond that the relevant difference between the two cases can be recognized via an appeal to the *Rule of Conservatism* (i.e. We may (very defeasibly) ignore those possibilities that those around us ignore). Lewis could rightly point out that those around us do (under normal circumstances) ignore the possibility that we acquired false beliefs because of typos in a paper. Sadly this will not solve the lottery case. The *Rule of Conservatism* says that we need not pay attention to the things other people ignore. It does *not* say that we must pay attention to the things around us that other people pay attention to. Susie is not required to make it salient to herself that she might win the lottery simply because it’s salient to other people.<sup>31</sup> More importantly, even if one did want to require Susie to do this by adding a sixth contextualist rule, Susie could still know that she would lose the lottery if Susie lived in a world where it was never salient to people that they might win the lottery. I take it that most people’s intuition is that no one can know that they will lose a fair lottery, and whether the possibility of winning is salient to the majority does not affect the strength of this intuition. So, adding a sixth contextualist rule will not work. At the same time, I have not (and cannot) consider all the possible rules that a contextualist might add to his account. So, the argument I am making against contextualism is not decisive. There could be some relevant difference between all lottery and all non-lottery cases that a contextualist might be able to discover. But so long as such a rule remains undiscovered, my argument hopefully retains its force. We are left with an account of the contextualist position which seemingly entails that we can know that we will lose some fair lotteries under certain conditions. This is a real

---

<sup>30</sup> So does Cohen’s fallibilism and (I would argue) all forms of contextualism.

<sup>31</sup> Or, if she is, it’s not because of what is required by the *Rule of Conservatism*.

problem since it contradicts the commonsense intuition that, as Lewis writes, “you do not know that you will lose the lottery, no matter what the odds are against you.”<sup>32</sup> It looks as though contextualists cannot have their cake and eat it too.

## VI. What Contextualists Ought to Say about the Lottery Paradox

If my argument works, then contextualists either face endorsing a view that entails near universal skepticism or endorsing a view that allows us to know that we can lose a fair lottery. Neither option seems particularly palatable. Does this mean that we ought to write contextualism off as a viable epistemological theory? Although I don't see how the contextualist can get around the two-horned dilemma I raise for the lottery paradox, I want to argue that contextualism can accept one of the horns relatively unscathed. The best defense for the contextualist in this case is a good offense. The lottery paradox is a problem for any epistemological view that allows for people to know some propositions are true when there is a non-zero chance that their belief is false. Except for skepticism, this seems to be a problem for every epistemological theory. Here is why. For any epistemological theory that allows someone to know a proposition is true when there is some non-zero chance the person could be mistaken, we can construct a fair lottery case where the odds of winning the lottery are exactly the same as the odds that the proposition one 'knows' is true is actually false. Given the intuition that most people have about lottery cases,<sup>33</sup> the result would be a tension in the epistemological theory in question. Either we know we will lose the lottery or we don't know the proposition in question or there is an inconsistency<sup>34</sup> in the view. It often seems to be an implicit assumption in the literature that the best way to account for our seemingly conflicting intuitions is to find the relevant difference between lottery cases and other cases of knowledge. This way we can show that our intuitions are really accurate and that we use the term 'knowledge' consistently. Given the vast literature in the field and, in my opinion, the lack of a satisfactory answer<sup>35</sup> to lottery cases, it is worth taking seriously the idea that we just use the term 'knowledge' inconsistently. Once this is accepted, we could revise the way we use the term to get out of the paradox. This is where contextualists can come out on top. Even though contextualism cannot sharply

---

<sup>32</sup> Lewis, “Elusive Knowledge,” 276.

<sup>33</sup> We can never know that we will lose a fair lottery.

<sup>34</sup> Or a consistent, but *ad hoc*, distinction.

<sup>35</sup> Even if the reader finds an account of lottery cases compelling, it should still be admitted that no explanation has achieved anything akin to widespread acceptance. Then again, that is standard in philosophical discourse.

distinguish all lottery cases from all non-lottery ones, Lewis' five rules get it right *most* of the time. Endorsing Lewis' contextualist rules in addition to the aforementioned definition of knowledge gets us as close as we can conceivably get to solving the problem, or so contextualists might want to argue. They are left with a much smaller bullet to bite (i.e. we can know that we can lose a fair lottery sometimes, in certain cases) than any form of invariantism.<sup>36</sup> In short, contextualists should admit that they cannot solve the lottery paradox, but then argue that no one else can either. At this point, the contextualist is in a good position to demonstrate that contextualist theories have the best resources to come closer to solving the paradox than any plausible alternative epistemic account. This actually provides us with a good reason to adopt a contextualist view.

## VII. Another Lottery Paradox: A Short Addendum

There is another lottery paradox, which Dana Nelkin refers to as the "rationality version," that is closely related to the one addressed in this paper.<sup>37</sup> What we want to say about this formulation depends upon what, exactly, we want to say about the first type of lottery case. I will define a few relevant terms, provide a brief explication of the problem and outline my solution. My solution is not definitive, but I hope that it contributes to the discussion and offers a new way of solving a second lottery paradox. Consider the following two plausible principles ...

1) *Sufficiency Thesis* (ST) "A proposition  $w$  is rationally acceptable if  $\text{Pr}(w) > t$ ."<sup>38</sup>

'Pr' is a probability distribution over propositions, while 't' represents a threshold value close to 1. The threshold can vary depending on the epistemological theory in question. But any non-skeptical view will require that 't' be less than 1.

2) *Conjunction Principle* (CP) "If each of the propositions  $w$  and  $c$  is rationally acceptable, so is  $w \wedge c$ ."<sup>39</sup>

---

<sup>36</sup> To be clear, this is because invariantist theories cannot account for the contextual differences between lottery cases (e.g. the possibility of winning is much more likely to be salient to the lottery player than possibilities of error are likely to be salient to the knowledge ascriber). So, the number of fair lottery cases that we could know we would lose should be greater for any invariantist account than it would for any contextualist theory.

<sup>37</sup> See Dana K. Nelkin, "The Lottery Paradox, Knowledge, and Rationality," *Philosophical Review* 109 (2000): 375-376, for an interesting discussion on the topic. Nelkin's solution is different from mine.

<sup>38</sup> Douven and Williamson, "Generalizing the Lottery Paradox," 755.

<sup>39</sup> Douven and Williamson, "Generalizing the Lottery Paradox," 755.

Given the acceptance of ST and CP, we get what John Hawthorne calls “the threat from conjunction introduction.”<sup>40</sup> He claims that an unpalatable consequence of the initial theory of relevant alternatives not only allows someone to know that they will lose the lottery, but also entails that they are able to predict everyone who will lose. For example, suppose that Johnny has entered a fair lottery with 5,001 ticket holders. He reasons that he will not win because the chance that this will occur is small enough to rule out. Furthermore, his friend Billy also has a lottery ticket. So Johnny can know that Billy will lose based upon the same reasons he knows that he will lose. Hawthorne argues that if Johnny knows that he will lose and Billy knows that Billy will lose, then Johnny can know that both he and Billy will lose. This becomes a problem if we suppose that Johnny knows 5,000 of the 5,001 ticket holders, in which he will safely be able to assert that he knows each one of them will lose. If we let the threshold in ST be less than 1 and accept CP, then the result is that people like Susie can predict who will win the lottery, which is absurd. To see why, consider the following case.

**The Disinterested Lottery Players** – Suppose Susie has eight siblings who share her feelings about the lottery and the moral nature of promises. Suppose that they each made the same promise to their mother that Susie did and play the lottery for the same reason. Lewis’ contextualism would entail that Susie can know each of her siblings, considered jointly, will lose.

For the sake of simplicity, let’s suppose the right threshold in ST is .8.<sup>41</sup> Also, let’s suppose that Susie and her siblings are playing a lottery with only 10 possible outcomes. Each person holds one ticket.

Now, given the conjunction principle, Susie could reason as follows.

1. My odds of winning are 1/10. So by ST I know that I will lose.<sup>42</sup> Call this knowledge claim [L]S
2. My first sibling’s odds of winning are 1/10. So by ST, I know S1 will lose. Call this knowledge claim [L]S1
3. My second sibling’s odds of winning are 1/10. So by ST, I know S2 will lose. Call this knowledge claim [L]S2
4. [Repeat for Siblings 3-8]

---

<sup>40</sup> Hawthorne, *Knowledge and Lotteries*, 94.

<sup>41</sup> This is, of course, way too low.

<sup>42</sup> Given the salience qualifier in the *Rule of Resemblance*, I grant Susie would not know this (on Lewis’ account) when she is thinking this issue through. But, presumably, she could know this on Lewis’ account as soon as she stopped thinking about the odds ceased to be salient in her mind.

Once Susie has reasoned her way through each of her siblings' cases, by CP, she can infer that none of her siblings will win the lottery even though there is a 90% that one of them wins. That is, since PR entails that Susie can know [L]S and she can know [L]S1, CP entails that Susie can know  $([L]S \wedge [L]S1)$ . She could eventually predict that there will be no winner (if the 10<sup>th</sup> option was not a ticketholder) or she could predict who the winner is (if the 10<sup>th</sup> option was a ticket holder). Obviously, Susie cannot do either one, and this needs to be accounted for by any plausible epistemological theory.

The line of thought that I think is worth developing rejects CP. If we (1) accept justification as a necessary condition for knowledge and (2) hold that justification depends on the subjective probability PR being sufficiently high, then (1) and (2) should entail that CP is false. Here is an illustration of the point.

**What Susie Can Know** – Everything true in the previous case is true here. Recall that it was stipulated that  $t = .8$  and the lottery consisted of 10 number variations. Given this, Susie can only know that, at most, two people who will lose at any given time. If we allowed CP, then Susie could come to know that large groups of people will lose, which would violate ST. That is, if we granted that Susie could know  $([L]S \wedge [L]S1) \wedge [L]S3$ , then  $PR = .7$ , which is below the set threshold. Anytime we allow CP and ST, agents can acquire knowledge about sets of things that are above  $t$  in ST.<sup>43</sup> The fact that accepting both CP and ST results in this tension means that we ought to reject CP.<sup>44</sup> In the case I considered, there are 55 different propositions Susie could know, but each are mutually exclusive. For example, Susie could know  $([L]S \wedge [L]S1) \text{ OR } ([L]S \wedge [L]S2) \text{ OR } ([L]S3 \wedge [L]S9)$ , but not more than one of these mutually exclusive combinations.<sup>45</sup>

Rejecting CP may be a hard sell, but it is arguably not as problematic as it may initially seem. The odds of Susie *and* S1 losing are greater than the odds of Susie losing *or* S1 losing. If we accept ST (as we should) and the idea that subjective probability is closely related to justification (as I think we should), then with each conjunction of lottery propositions (via CP), a certain amount of justification is lost. By the time you have a conjunction of knowledge propositions about every ticket holder, justification has withered away completely.<sup>46</sup>

---

<sup>43</sup> And this will be true no matter what  $t$  is set as.

<sup>44</sup> The rejection of PR would result in skepticism, which we want to avoid if at all possible.

<sup>45</sup> The 'OR' in this sentence should be read as the exclusive, not inclusive, disjunction.

<sup>46</sup> Hawthorne has written a fair amount on the issue and there is some interesting discussion in John Hawthorne and Luc Bovens, "The Preface, the Lottery, and the Logic of Belief," *Mind* 108 (1999): 241-264.

### VIII. Conclusion

After providing a short overview of Lewis' contextualism and how it is meant to handle lottery cases, I raised a few unaddressed issues about lotteries. The main argument I advance is that Lewis' contextualism is either committed to saying that we can know that we will lose a fair lottery or adopting something close to universal skepticism. Lewis does not consider lottery cases where there is no guaranteed winner. I argue that there is no relevant difference between lotteries where there is a guaranteed winner and where there is no guaranteed winner. Because each case is relevantly similar, we ought to treat them as such.<sup>47</sup> I then offered a thought experiment to demonstrate that Lewis' contextualism allows for a person to know she can lose the lottery, as none of his contextualist rules are violated in my case. Next, I argued that the lottery cases generalize and that Lewis' contextualism could not deny knowledge in these lottery cases without also having to deny knowledge in every case where we could be mistaken. Since skepticism would be worse than a few counter-intuitive knowledge claims, I concluded that contextualists should bite the bullet and allow that we can know we will lose a fair lottery sometimes. On contextualists' behalf, I argued that they are in a better position to handle the lottery paradox than any invariantist account. I end the paper by considering another lottery paradox that contextualists who adopt my proposed solution will inherit and argue that they can get out of this lottery paradox by denying the conjunction principle.<sup>48</sup>

---

<sup>47</sup> Even if both types of lottery cases are not treated the same, Lewis' Contextualism would not be off the hook. It still could not account for our intuitions about those lottery cases with no guaranteed winner.

<sup>48</sup> I would like to thank André Gallois, Matt Eller and Amanda O'Neil for helpful comments and feedback on earlier drafts of this paper.