

HOW TO UNDERSTAND AND SOLVE THE LOTTERY PARADOX

Patrick BONDY

ABSTRACT: It has been claimed that there is a lottery paradox for justification and an analogous paradox for knowledge, and that these two paradoxes should have a common solution. I argue that there is in fact no lottery paradox for knowledge, since that version of the paradox has a demonstrably false premise. The solution to the justification paradox is to deny closure of justification under conjunction. I present a principle which allows us to deny closure of justification under conjunction in certain kinds of cases, but which still allows that belief in a conjunction on the basis of justified belief in its conjuncts can often be justified.

KEYWORDS: lottery paradox, knowledge, justification, closure

The purpose of this paper is to explain the correct way to understand the lottery paradox, and to show how to resolve it. Briefly, the lottery paradox goes as follows. In a fair lottery, there is a high probability that any given ticket will lose (say, 0.999, for a 1000-ticket lottery), and the same goes for every other ticket. If you buy a ticket in a fair lottery, you would therefore be justified in believing that your ticket is a loser, and you would be similarly justified in forming that belief of each other ticket as well. You would therefore also be justified in believing of *all* of the tickets that they will lose. But you are also justified in believing that one ticket will win, since you know that it is a fair lottery. So you are justified in believing both a proposition (all of the tickets will lose) and its negation (at least one ticket will not lose). And that certainly looks paradoxical.

This paper is divided into two parts. In the first, I explain how to properly understand the lottery paradox. In particular, I argue that although there does appear to be a paradox that needs to be resolved when we cast the problem in terms of justification, there is no paradox when we cast the problem in terms of knowledge. Contra Dana Nelkin¹ and Jonathan Sutton,² then, it is not a mark in favour of a solution to one formulation of the paradox, that it also offers us a solution to the other. In the second part of the paper, I argue that the correct

¹ Dana Nelkin, "The Lottery Paradox, Knowledge, and Rationality," *Philosophical Review* 109 (2000): 373–409.

² Jonathan Sutton, *Without Justification* (Cambridge, Mass: MIT Press, 2007).

Patrick Bondy

solution to the lottery paradox is to deny the closure of justification under conjunction.

1. How to understand the lottery paradox

Jonathan Sutton,³ following Dana Nelkin,⁴ sets out two versions of the lottery paradox, one for knowledge and one for justification. I will argue that the justification version is a serious problem which requires a solution, but the knowledge version is not a problem at all. The two versions are as follows:

The Knowledge Paradox

1. Jim knows that his ticket t_1 will lose.
2. If Jim knows that his ticket t_1 will lose, then he knows that t_2 will lose, he knows that t_3 will lose, ... and he knows that $t_{1,000,000}$ will lose.

So,

3. Jim knows that t_1 will lose ... and Jim knows that $t_{1,000,000}$ will lose. (1,2)
4. Jim knows that either t_1 will not lose or t_2 will not lose ... or $t_{1,000,000}$ will not lose.
5. Propositions of the following form comprise an inconsistent set: (a) $p_1 \dots (n) p_n$, $(n+1)$ not p_1 or ... not $p(n)$.

So,

6. Jim knows propositions that form an inconsistent set. (3,4,5)
7. It is not possible to know propositions that form an inconsistent set.

So,

8. (1), (2), (4), (5), or (7) is false.

The Justification Paradox

- 1*. Jim could justifiably believe that his ticket t_1 will lose.
- 2*. If Jim could justifiably believe that his ticket t_1 will lose, then he could justifiably believe that t_2 will lose, he could justifiably believe that t_3 will lose ... and he could justifiably believe that $t_{1,000,000}$ will lose.

So,

³ Sutton, *Without Justification*, 49-50.

⁴ Nelkin, "The Lottery Paradox."

How to Understand and Solve the Lottery Paradox

3*. Jim could justifiably believe that t_1 will lose ... and Jim could justifiably believe that $t_{1,000,000}$ will lose. (1*, 2*).

4*. Jim could justifiably believe that either t_1 will not lose or t_2 will not lose ... or $t_{1,000,000}$ will not lose.

5*. Propositions of the following form comprise an inconsistent set: (a) $p_1 \dots (n) p_n, (n+1)$ not p_1 or ... not p_n .

6*. Jim recognizes that the following propositions form an inconsistent set: (i) t_1 will lose ... (n) $t_{1,000,000}$ will lose, (n+1) either t_1 will not lose ... or $t_{1,000,000}$ will not lose.

So,

7*. Jim could justifiably believe inconsistent things that he recognizes are inconsistent. (3*, 4*, 5*, 6*)

8*. One cannot justifiably believe things that one recognizes are inconsistent.

So,

9*. (1*), (2*), (4*), (5*), (6*), or (8*) is false.

Both Nelkin and Sutton take it to be best if a theory is able to give parallel solutions to the two versions of the paradox. However, I will argue that it is quite clear that the two versions must be given different treatments.

The justification paradox appears to present a serious difficulty for theories of justification. One of the main solutions proposed in the literature is to deny premise (1*), and hold that however probable Jim's belief that his ticket will lose may be, he is not justified in believing it. Simon Evnine,⁵ for example, defends this solution, on the grounds that beliefs that are members of "Indifferent Sets" – sets of beliefs one of which must be false but none of which has anything to recommend it over any other – are not rational to believe.

A second proposed solution is to deny premise (2*) and hold that, even if Jim is justified in believing of ticket t_1 that it will not win, he is not similarly justified with respect to (some of the) remaining tickets. Gilbert Harman⁶ proposes such a solution. He argues that we are justified in believing that the first ticket will lose because the odds against its winning are 999,999 to 1. Likewise, we are justified in believing that the second ticket will lose. However, we are not justified in quite the same way for this second belief as we are for the first. For the second belief, we are justified in our belief that the ticket will lose because the odds against it are 999,998 to 1, rather than 999,999 to 1. This is because we must take

⁵ Simon Evnine, "Believing Conjunctions," *Synthese* 118 (1999): 201-227.

⁶ Gilbert Harman, *Change in View* (Cambridge: MIT Press. 1986).

account of our prior justified belief that ticket n will lose, which decreases the effective size of the lottery. This process continues for the beliefs we generate for each ticket in turn, and at some point in this process the odds become too uncertain to justify the belief of each remaining ticket that it will lose.

Another solution is to deny premise (8*), and claim that it is possible to justifiably believe things one knows to be inconsistent. Of course, defenders of this third solution have to come up with some weaker form of premise (8*) in order to avoid the lottery paradox without licensing wholesale inconsistency in our beliefs. Nelkin considers and rejects one such weaker principle, which she calls the Foley Principle (FP): it cannot be rational to believe a proposition that is internally inconsistent.⁷ An example of an internally inconsistent belief would be a conjunction with conjuncts that cannot be jointly true. Thus, one might be justified in each of the individual beliefs that makes up the lottery paradox, but not justified in believing the (internally inconsistent) conjunction of these beliefs. This solution to the paradox amounts to a rejection of the closure of justification under conjunction.

None of these proposals is very tempting, on the face of them. In the version of the paradox discussed by Nelkin and Sutton, the odds that Jim's belief is true are 999,999 to 1; it is hard to imagine an empirical belief that one could be better justified in believing. The denial of premise (1*) therefore appears to lead directly to scepticism about empirical justification for our beliefs. The denial of premise (2*) also seems problematic. If we accept Harman's argument, then it seems that Jim learns new information about the lottery based on the order in which he forms his beliefs. For example, if he starts at ticket n , then it seems as if Jim can deduce that he ought to buy a ticket from the second half of the lottery, one of tickets £500,000–£1,000,000, since then he will be more likely to win. If Jim started his considerations on the odds of each ticket winning from ticket £1,000,000, however, he would be justified in coming to exactly the opposite conclusion. This apparently absurd result illustrates the difficulty with denying premise (2*). Finally, to deny premise (8*) and replace it with some weaker principle such as (FP) is also apparently a hard pill to swallow: a conjunction is true just in case each of its conjuncts is true, so if Jim is justified in believing each of the conjuncts, he ought also to be justified in believing the conjunction. Rejecting (8*) puts us in the bizarre position of being blocked from performing normally innocuous logical operations on our body of justified beliefs. This is counterintuitive on the face of

⁷ Richard Foley argues in various places that a principle like this one ought to replace one like (8*). See, for example, his "Justified Inconsistent Beliefs," *American Philosophical Quarterly* 16 (1979): 247–257.

it, and presents the challenge of finding a principled distinction between cases where the conjunction of justified beliefs is permitted and cases where it is blocked.

The point here is not that none of these proposals can work; it is just that none of them are initially very appealing, so it is not immediately obvious how to respond to the justification version of the lottery paradox. However, a complete theory of epistemic justification that does not embrace scepticism with respect to empirical justification will have to adopt a solution along one of these lines.

The knowledge version of the paradox, on the other hand, has a clear solution; in fact, it is not really a paradox, because premise (2) is demonstrably false. Before going into the demonstration, though, notice that even if (2) were not demonstrably false, there would be an intuitively plausible solution to the puzzle, in the denial of (1). Many people have the intuition that one cannot know that one's lottery ticket is a loser before the winner has been drawn, so this solution is likely not a difficult one to sell. Furthermore, rejecting premise (1) provides a good explanation for why people ever buy lottery tickets at all. If people know in advance that their tickets will be losers, then the phenomenon of lottery-ticket-buying calls for explanation. Another reason to deny premise (1) might be that Jim has not ruled out, and cannot rule out, the relevant alternative that his ticket is a winner. That alternative is relevant here, because it is a very close possible world in which Jim wins the lottery – all that has to happen for that possibility to be actualized is that his ticket be drawn in the lottery. On top of those reasons, if the knowledge paradox was a real paradox, we would have yet another reason to reject premise (1), and deny knowledge in lottery cases.

However, we do not need to deny premise (1) in order to escape the paradox. We might want to reject it for those other reasons, but the knowledge paradox gives us no reason to do so, because premise (2) is demonstrably false. Premise (2), again, is that if Jim knows that his ticket t_1 is a loser, then he knows that t_2 is a loser, ... and he knows that $t_{1,000,000}$ is a loser. However, even if Jim knows that his ticket t_1 is a loser, it does not follow that he knows each of " t_2 is a loser," " t_3 is a loser," ... and " $t_{1,000,000}$ is a loser," even though he will have the same degree of epistemic justification for each of them, because one of them will be false. One of the tickets is, or will be, a winner, and Jim cannot know of *that* ticket that it will lose, since one cannot know a falsehood. Therefore, premise (2) is clearly false, for even if Jim does know that ticket t_1 will lose, he cannot know this of all the other tickets, but (at most) of all but one of the other tickets.⁸

⁸ A possible objection here is that premise (2) is a conditional, so it is true if its antecedent is false. Since I claim in this paper that it is possible that premise (1) is false, and Jim does not

This result will hold for any definition of knowledge that includes truth as one of its constituents. Since both Sutton and Nelkin accept that truth is part of knowledge, each will have to accept that premise (2) is false, for reasons independent of any attempt to resolve the supposed paradox. Without premise (2), the paradox cannot get going. In order to generate the paradox, Jim is required to know an inconsistent set of propositions, as stated in premise (6). But there is nothing inconsistent about knowing of all but one ticket that each will not win, and knowing that exactly one ticket will win. This is perfectly consistent, and just what we would expect. The knowledge paradox is therefore not a paradox at all.

There is no analogous solution to the justification paradox. Although the solution of the knowledge paradox is to deny premise (2), and one of the available solutions to the justification paradox is to deny the analogous premise (2*), we are not in a position to give a demonstration of the falsity of (2*) as we are of (2). The reason for wanting to deny (2*) in the justification paradox is just the desire to escape the paradox without rejecting premise (1*) or (8*). The reason for rejecting premise (2) in the knowledge paradox is that it can be independently demonstrated to be false: Jim cannot know of each ticket that it will lose, because one of them is a winner. The solution to the knowledge paradox therefore does not give us any indication about how we ought to try to resolve the justification paradox.

This should not be a surprising result. The justification paradox arises because Jim seems to be in an identical epistemic position with respect to each of the lottery tickets. Whatever he is justified in believing about one ticket, there seem to be no non-arbitrary grounds for denying that he is justified in believing precisely the same thing about every other ticket. This is why Harman's rejection of (2*) is so problematic; it allows Jim to arbitrarily treat some of the tickets differently from the others.

However, there is at least one sense in which the tickets are not all identical: one and only one ticket will win. This distinguishing feature is precisely

know that his ticket will lose (i.e. the antecedent of (2) is false), I should therefore hold that (2) could be true. There are two points to note here. First, this objection requires that we treat the ordinary-language conditional as the material implication of traditional logic, and that analysis of the conditional is by no means uncontroversial. Second, even granting that analysis of the conditional, my argument can be recast without affecting the main point, as follows: it is demonstrably true that either premise (1) is false or premise (2) is false. Although there are reasons to reject (1), I do not take a stand regarding (1). My claim is that *if*(1) is true, then (2) is false. Putting the argument this way does not change the fact that the knowledge paradox necessarily has a false premise, nor does it change the fact that the reason to reject premise (1) or (2) has nothing at all to do with the desideratum of resolving a paradox.

what is made relevant by the switch from justification to knowledge. Knowledge, because it is factive, is able to take account of the difference between the winning ticket and all other tickets. This breaks the symmetry that holds between the beliefs about t_1 , t_2 , etc., upon which the lottery paradox depends. Justification, on the other hand, cannot differentiate the winning ticket from the other tickets by the sheer fact that it will win.

Sutton uses the supposed analogy between the knowledge and justification versions of the lottery paradox to argue for his claim that justification just is knowledge.⁹ He endorses the rejection of premise (1) as a solution to the knowledge paradox, and then points out that his account of justification-as-knowledge entails the rejection of premise (1*) in the justification paradox as well. If justification is knowledge, and Jim does not know that his ticket is a loser, then he is not justified in believing that his ticket is a loser, either. Score one for the justification-as-knowledge thesis: it solves the two paradoxes in the same way.

A recent objection to this argument of Sutton's can be found in Coffman.¹⁰ Coffman's own solution to the knowledge paradox is to deny Jim's knowledge that his ticket is a loser (that is, to reject (1)), and his solution to the justification paradox is to reject the closure of justification under conjunction (in effect, although Coffman sets up the paradox in a slightly different way, this solution is to reject (8*)).

In favour of Sutton's solution to the paradoxes is that it is elegant, since the two paradoxes are given a unified solution. Counting against it is the fact that it involves the counterintuitive denial of the justification of Jim's belief that his ticket will lose. Coffman's solution, on the other hand, is piecemeal, but it respects the intuition that Jim's belief that his ticket is a loser is justified. Coffman argues that being piecemeal is not a significant mark against a solution to the paradoxes, so his solution is at least as plausible as Sutton's.

Coffman's argument is fine as far as it goes, but it does not go far enough. He has argued that there is another solution to the two versions of the paradox that is at least as plausible as Sutton's, so the fact that Sutton's view of justification-as-knowledge gives a unified solution to the paradoxes does not count as a reason to accept that view. But what I have argued here goes much further than that: the way to deal with the knowledge paradox is to point out that it is not a paradox at all, because it has a clearly false premise. It is therefore not even a desideratum that a theory be able to offer a unified solution to the paradoxes.

⁹ Sutton, *Without Justification*, 51.

¹⁰ E. J. Coffman, "Is Justified Belief Knowledge? Critical Notice of Jonathan Sutton, *Without Justification*," *Philosophical Books* 51 (2010): 16-17.

2. How to solve the lottery paradox

The only way to formulate the lottery paradox, then, is in terms of justification, not knowledge. In fact, it seems to me that there is a perfectly good solution to the justification version of the paradox, but it does not mirror the solution to the knowledge version. The solution is to deny the closure of justification under conjunction, which, as I have said, amounts to a rejection of (8*). Coffman offers this same solution, but he only puts it forward as an equally plausible solution as Sutton's denial of premise (1*), without offering an argument in support of it. Richard Foley¹¹ also argues for denying the closure of justification under conjunction, but his strategy is to argue that closure of justification under conjunction has absurd consequences, and to argue that other solutions to the lottery paradox are simply worse.

Here, then, is the positive argument for this solution to the lottery paradox. To keep things simple in what follows, I reduce the number of tickets in the lottery to 100, but the point remains the same for a lottery of 1,000,000 tickets.

When Jim considers whether a given ticket t_1 is a loser, he knows that there is a probability of 0.99 that it is, since there are ninety-nine equally probable ways for t_1 ticket to lose, and only one way for it to win. He is therefore justified in believing that that ticket is a loser. He also knows that each other ticket has the very same probability of winning. He is therefore justified in believing of each one that it will lose. But when Jim considers whether *both* t_1 *and* t_2 will lose, we can see that there are ninety-eight ways for that proposition to turn out true (i.e. it is true just in case t_3 wins, or t_4 wins, ... or t_{100} wins), and only two ways for it to turn out false (if either t_1 wins or t_2 wins). So the probability that the proposition that tickets t_1 and t_2 are losers is true is 98/100, or 0.98. And so on: as we increase the size of the set of tickets that we believe are losers, we lower the probability that our belief is true. Once we reach the end of the tickets, and we consider whether t_1 will lose, *and* t_2 will lose, *and* t_3 will lose... *and* t_{100} will lose, it is obvious that there are 0 ways for that conjunction to come out true. So there is no justification whatsoever for believing the conjunction that all of the lottery tickets will lose.

The solution to the lottery paradox, then, is to allow that there is very good justification for believing of each individual ticket that it will lose, and to allow that there is very good (albeit slightly less) justification for believing that a given set of two tickets will lose, but to insist that there comes a set of tickets that is

¹¹ Foley, "Justified Inconsistent Beliefs."

sufficiently large that the probability that they are all losers is sufficiently low that the belief that they are all losers is not justified.

Trying to point out the exact point where a set of tickets is sufficiently large that we no longer have justification for believing that all of the tickets in that set are losers is rather like trying to point out the point at which a man goes from not bald to bald. But the difficulty of identifying that point is no reason to doubt that there is a boundary (perhaps a vague one) between the two types of case. Just as there are clear cases where a man is bald, and clear cases where he is not, there are also clear cases where the belief that a set of tickets are all losers is justified, and clear cases where such a belief is not.

Can this type of solution to the lottery paradox can be made to work in a straightforward way for other similar paradoxes, such as the preface paradox? Perhaps, but I am skeptical. Briefly, the preface paradox asks us to consider an author of a book who knows that she has done her research well, she is a careful writer, and so on, but she is not so bold as to believe that every single statement in the book is true. She therefore writes a modest preface in which she claims to be sure that she must have made at least a few mistakes. Nevertheless, she still believes, of each proposition in the book, that it is true.

The preface paradox is clearly structurally similar to the lottery paradox. In each case, we have an agent who holds a set of beliefs each of which is well justified, but who does not believe that the conjunction of those beliefs is true. It would therefore be nice to have a similar solution to both paradoxes. Perhaps a solution of the sort that I have offered to the lottery paradox can be made to work for the preface paradox. However, I doubt that it can be applied in a straightforward fashion, because in the preface paradox, the probabilities of each of the propositions in the book are not clear. We cannot simply count up the number of ways to be mistaken and the ways to be correct, and yield a definite judgment about the probability that a given set of propositions is true.

Still, it does at least look like denying closure of justification under conjunction is the way to solve the lottery paradox. As I point out above, however, denying closure in this way makes it incumbent upon me to put forward a way to block justified conjunction in the lottery case, while allowing it in ordinary cases that do not appear to be problematic. The principle that I propose is this:

Improbable Conjunctions (IC)

In cases where the justification of a belief is determined by its probability, and conjoining two or more beliefs that are independently probable yields a

conjunction which is sufficiently less probable than the belief in its conjuncts, the belief in that conjunction on the basis of its conjuncts is not justified.

This principle is different from the Foley Principle (FP) which Nelkin considers. FP says that it cannot be rational to believe internally inconsistent propositions. FP blocks justified inference to the belief: “one lottery ticket is a winner and all of the lottery tickets are losers,” because that proposition is internally inconsistent, and so FP blocks the lottery paradox in its normal form. But FP does not block the inference to the belief: “one ticket is a winner and each of t_1 through t_{n-1} is a loser,” in an n -ticket lottery, which it seems to me we should also want to say is an unjustified belief. By contrast, the principle IC blocks the inference to that further belief as well, since it is a highly improbable belief, despite the fact that it is a conjunction and each of its conjuncts is independently very probable.

Principle IC also allows that in a wide range of ordinary cases where we believe conjunctions on the basis of belief in the conjuncts, the resulting belief will be justified, which is another very important desideratum. IC allows such inferences in cases where a conjunction of probable propositions is itself still sufficiently probable. The principle admittedly does not specify a point at which the probability of a conjunction becomes too low for justified belief in it – but then, neither do typical accounts of justification specify probabilities for justified and unjustified beliefs. Specifying the threshold of probability for justification is not a requirement for holding that there is such a threshold, even if only a vague one.

Indeed, on reflection, it should not even be surprising that we have to deny closure of justification under conjunction. If it is because a belief is sufficiently probable that it counts as justified, as in the case of the lottery-beliefs, then of course there will be cases where belief in a conjunction is unjustified even though the belief in the conjuncts is justified, since a conjunction is usually less probable than its conjuncts taken independently. Sometimes a conjunction will be highly improbable, despite having independently plausible conjuncts. Now, I do not claim that it is always a belief's probability that determines its justification, but in the case of lottery beliefs, it makes good sense to think so, and in such cases, it is clear that we must deny closure of justification under conjunction.¹²

¹² I am very grateful to Benjamin Wald, whose ideas and feedback have been central to the development of this paper.