

Four Arguments for Denying that Lottery Beliefs are Justified¹

Martin Smith

A 'lottery belief' is a belief that a particular ticket has lost a large, fair lottery, based on nothing more than the odds against it winning. The lottery paradox brings out a tension between the idea that lottery beliefs are justified and the idea that that one can always justifiably believe the deductive consequences of things that one justifiably believes – what is sometimes called the *principle of closure*. Many philosophers have treated the lottery paradox as an *argument* against the second idea – but I make a case here that it is the first idea that should be given up. As I shall show, there are a number of independent arguments for denying that lottery beliefs are justified.

Keywords: Lottery belief, lottery paradox, closure, justification

I. Introduction

Suppose I hold a ticket – ticket #5472 say – in a fair 10000 ticket lottery with a single winner. Suppose the lottery has been drawn, but I'm yet to hear the result. Suppose I already believe, however, that ticket #5472 has lost, based on the fact that there is only one winning ticket and 9999 losers. Surely, given these odds, I am justified in believing that ticket #5472 has lost. There's nothing special, however, about the ticket that I happen to be holding – it has as good a chance of being the winner as any other ticket. As such, if I'm justified in believing that ticket #5472 has lost, then I should be justified in believing the same thing about ticket #1, about ticket #2, about ticket #3 ... right up to ticket #10000.

If these 10000 'lottery beliefs' were all justified – if I justifiably believed that ticket #1 has lost, justifiably believed that ticket #2 has lost, justifiably believed that ticket #3 has lost etc. then, given that I know there are only 10000 tickets in the lottery, I could simply *deduce* that every ticket has lost. The deduction is very simple, so surely I would be justified in drawing this conclusion. And yet, I also know that one of the tickets must have won. From these beliefs I can now deduce a contradiction: Every ticket has lost and one of the tickets has won. Clearly, though, I could not justifiably believe *this*. This is the lottery paradox – a simple piece of reasoning in which one can be easily caught up and swept along, but which leads us to an absurd conclusion.

Many epistemologists grant that I can justifiably believe that ticket #5472 has lost, justifiably believe that ticket #1 has lost, justifiably believe that ticket #2 has lost etc. but think that the reasoning goes wrong at the next step – for I would not be justified in putting these beliefs together and concluding that every ticket has lost (Kyburg, 1961, 1970, Schick, 1966, Derksen, 1978, Foley, 1993, chap. 4, 2009). Those who take this line are forced to reject what is sometimes called the

¹ This paper was presented at the University of Glasgow in November 2015. Thanks to everyone who participated on this occasion. Particular thanks to Campbell Brown and Stephan Leuenberger for a number of very helpful discussions of this material.

principle of *closure*, according to which deductive inferences always preserve justification. This might be set out a bit more precisely as follows:

Closure If one justifiably believes P_1 , justifiably believes P_2 , justifiably believes P_3
 ... and comes to believe Q by competently deducing it from P_1, P_2, P_3 ...
 then one justifiably believes Q .

The deductive inference involved in the lottery paradox, in effect, involves nothing more than *conjoining* a number of beliefs. As such, those who would solve the lottery paradox by denying closure are obliged to also deny a particular instance of the principle, concerning inferences in which a number of premises are simply conjoined.

The closure principle has some intuitive appeal. It's natural to think that competent deduction should be an epistemically safe way of expanding one's body of beliefs, and should never take one from the realm of justified belief to the realm of irresponsible speculation. For many epistemologists, though, the lottery paradox provides a powerful argument against this principle. Other epistemologists, emphasising the plausibility of closure, have treated the lottery paradox as an argument against the claim that lottery beliefs are justified – an argument for denying that I can justifiably believe, of any particular ticket, that it has lost, based purely on the odds involved (Lehrer, 1974, chaps. 6, 8, Ryan, 1996, Douven, 2002)².

My sympathies lie with the second group – I would also be inclined to deny that lottery beliefs are justified. And yet, to deny this purely, or even partly, out of a desire to preserve closure is somewhat dissatisfying – and leads to an immediate deadlock with the closure-deniers. In this paper, I will attempt to avoid this dialectical impasse by outlining a number of independent arguments for denying that lottery beliefs are justified. In the coming sections, I will outline four such arguments, each of which makes use of certain general suppositions about justification, but none of which makes use of the closure principle. Undoubtedly there are certain connections between the arguments I will offer, and the kinds of sentiments that they appeal to, but I present them here as separate. Any one of these arguments could be accepted independently of the others – and, in order for this paper to achieve its aim, it is enough that just one of them be accepted.

Before moving on it may be worth saying something, briefly, about *subjective probability*. I have heard it claimed that the lottery paradox is no 'real' paradox at all, because once we start thinking in terms of subjective probabilities, rather than beliefs, it's obvious what I should do when confronted with the lottery situation: When it comes to the proposition that ticket #5472 has lost, I should assign it a subjective probability of 0.9999. The same goes for the proposition that ticket #1 has lost, the proposition that ticket #2 has lost, the proposition that ticket #3 has lost etc. The proposition that ticket #1 has lost and ticket #2 has lost should get a subjective probability of 0.9998,

² A third group have suggested that one can justifiably believe, of *some* lottery tickets, that they have lost, but one cannot believe this about too many of the tickets at once (Harman, 1986, pp70-72, Kroedel, 2012, see also Douven, 2008). This solution to the lottery paradox allows us to retain closure *and* grant that lottery beliefs can be justified. But the solution carries significant costs of its own. It predicts, for instance, that the justificatory status of a new lottery belief will depend on what lottery beliefs one already holds – that one lottery belief may count as less justified than another simply because it was formed *later*, even though both beliefs are equally supported by one's evidence. I won't discuss this view any further here (for some discussion see Littlejohn, 2012).

the proposition that ticket #1 has lost and ticket #2 has lost and ticket #3 has lost should get a subjective probability of 0.9997 and so on, while the proposition that all the tickets have lost should get a subjective probability of 0.

In so far as I regard the lottery set-up itself as being certain, this does seem a very plausible description of how I should distribute my subjective probabilities and, sure enough, it does not lead to anything paradoxical. But demonstrating that we can give a consistent description of my epistemic obligations in terms of subjective probabilities doesn't *solve* anything exactly – it doesn't show us, for instance, where the reasoning in the lottery paradox goes wrong. In fact, as best I can tell, this description can be reconciled with any of the possible solutions to the lottery paradox. In particular, it is perfectly compatible with denying that lottery beliefs are justified. This is something I will periodically return to.

II. Justification and Knowledge

Even if ticket #5472 has in fact lost the lottery, in the situation described I don't yet *know* that it has lost the lottery. Perhaps I could come to know this when I hear the result announced, but I cannot know it purely on the basis of the odds involved. Lottery beliefs do not constitute knowledge. This is widely accepted (see, for instance, Harman, 1968, Dretske, 1971, DeRose, 1996, Nelkin, 2000, Hawthorne, 2003, partic. chapter 1, and Pritchard, 2007)³ and serves as the starting premise for the first argument. Some epistemologists have suggested that the only way a belief can be justified is if it constitutes knowledge (Sutton, 2007, Littlejohn, 2017, Williamson, forthcoming). If that's correct, then it will follow immediately that lottery beliefs are not justified and we're done⁴. For most epistemologists, though, this quick argument will seem suspect – most epistemologists allow that justified beliefs can fall short of knowledge.

Suppose I wander into a room for the first time, have a perceptual experience as of a red wall before me, and immediately come to believe that there is a red wall before me. This would seem a very clear case of a justified belief – and yet, for all that has been said so far, the belief may not be knowledge. It might be that the wall is actually white but illuminated by deceptive red light, emanating from a hidden source, or it might be that the wall is actually green but I've been suddenly struck by colour blindness or am undergoing a vivid colour hallucination. In these circumstances I would justifiably believe that the wall is red, but I wouldn't know that the wall is red. Perhaps lottery beliefs also belong in this kind of category – perhaps they are further examples of beliefs that are justified, but not knowledge. On closer inspection though, my red wall belief seems not at all like a lottery belief.

³ This judgment also appears to be widely accepted amongst non-philosophers – see Turri and Friedman, 2014, Friedman and Turri, 2015 and Ebert, Smith and Durbach, 2018.

⁴ One of Sutton's arguments for identifying knowledge and justified belief is that this allows us to deny that lottery beliefs are justified and to resolve the lottery paradox without abandoning closure (see Sutton, 2007, section 2.1.2). The argument, however, generates little force – even for those who are intent on maintaining closure. As I hope to show, the claim that lottery beliefs are not justified can be motivated in a range of less radical ways.

If my belief that there is a red wall before me fails to constitute knowledge, then there must be something unusual about the circumstances in which I find myself – either there's tricky lighting, or my perceptual system is malfunctioning etc. If the circumstances had been more conducive – if the wall had been red, and the surrounding environment and perceptual process had been normal – then my belief would have been knowledge. In contrast, when it comes to my belief that ticket #5472 has lost, its failure to constitute knowledge is not due to any unusual circumstances. Even if we make the circumstances as conducive as possible – the lottery is drawn as expected, everything is normal and above board, ticket #5472 really does lose etc. – *still* I don't know that ticket #5472 has lost.

Consider some further examples. Suppose I'm asked where my car keys are and I immediately reply that they're on the kitchen table, having a clear memory of putting them there just minutes ago. My belief concerning the whereabouts of my car keys clearly seems justified – but we could still imagine that it fails to be knowledge. We could imagine that the keys have been stolen in the last two minutes by a stealthy, fast-working burglar, or that I never even put my car keys on the table, but am afflicted with a false or fabricated memory. Whatever the case though, if my belief fails to be knowledge, this will be due to circumstances conspiring against me in some way.

Suppose, finally, that I'm told by a colleague that the staff meeting starts at 2pm and I take her at her word. Provided I have no reason to doubt my colleague's testimony, this belief would be justified and, provided the meeting does start at 2pm, and my colleague is just passing the information along in a straightforward way, the belief would be knowledge as well. If the belief is not knowledge, then this will be because my colleague means to deceive me, or has been deceived herself, or the meeting time has been changed at the last minute etc. If the belief fails to be knowledge, then the failure will be attributable to some factor of this kind.

These examples represent a fair cross-section of the kinds of evidence that can serve to justify beliefs. And, in each example, this evidence is also perfectly sufficient for knowledge. In each example, we can fill in the external, extra-evidential circumstances in such a way that the belief would be knowledge. My evidence for believing that ticket #5472 has lost seems very different – this evidence could never be enough for me to know that ticket #5472 has lost, irrespective of how the external circumstances are filled in.

If we held (with Sutton, Littlejohn and Williamson) that a belief must be knowledge in order to count as justified then, as noted at the beginning of the section, we would have immediate reason to deny that lottery beliefs are justified. But this denial can be motivated by a much weaker claim regarding the connection between justification and knowledge, and one that should be acceptable to a broad range of epistemologists: In order for a belief to count as justified it must, at least, be a *candidate* for knowledge – it must be a belief that could count as knowledge, given the right external conditions (see Bird, 2007, Littlejohn, 2012, Ichikawa, 2014, Smith, 2014, 2016, chap. 1).

This generalisation is strongly supported by the examples just considered. Some of the most clear-cut cases of justified beliefs, including beliefs justified by perceptual, testimonial and memorial evidence, are all clear candidates for knowledge. To insist that lottery beliefs are justified is not just to admit some new justified beliefs, it is to admit a whole new *category* of justified beliefs – to admit a kind of evidence that can make for justification but cannot make for knowledge. This, at the very least, is something that should give us serious pause.

If a belief must be a candidate for knowledge in order to count as justified, then I cannot justifiably believe that ticket #5472 has lost the lottery based purely on the odds against it winning. But what this principle *will* permit me to justifiably believe is that ticket #5472 is very likely to have lost, or that ticket #5472 is 99.99% likely to have lost. These *are* things that I can plausibly know, purely on the basis of the odds involved. By going further than this, and believing that ticket #5472 *has* lost, I overstep the bounds of potential knowledge and, on the present thinking, overstep the bounds of justification as well.

III Justification and Surprise

The second argument begins from the premise that, if I believe a proposition P, and my belief is justified, then the falsity of P should be *surprising* for me. While there is substantial disagreement amongst epistemologists as to the ultimate aims of belief, almost all would accept that one of our aims, as believers, is to avoid believing things that are false. If I believe a proposition P, even though the falsity of P should come as *no surprise* to me, then it's plausible that my belief *flouts this goal*. And if a belief flouts this goal then, whatever else we might want to say about it, we should not regard it as justified.

At first, it might seem as though this principle poses no threat to the justificatory status of lottery beliefs. Surely it *would* be surprising for me if my ticket, ticket #5472, won the lottery – wouldn't I gush about how 'surprised' I was if my ticket was actually declared the winner? The relevant question, however, is not whether I *would* feel surprised in the event of my ticket winning, but rather whether I *should* be surprised by such an event – and this, I think, is somewhat less clear. For the moment, instead of focussing on the ticket that I happen to own, let's consider a different ticket – ticket #6255 say. Suppose I have no idea who holds this ticket and no particular interest in it. Should I be surprised if ticket #6255 turned out to be the winner?

As a number of philosophers have observed, the mere fact that an event is infrequent or improbable is not enough for it to count as surprising (Hambourger, 1980, pp594-595, Horwich, 1982, chap. 5, Peacocke, 2004, pp84-85, Smith, 2017). Suppose I roll a 12-sided die, a 10-sided die and a 6-sided die. Suppose they come up 11, 7 and 4. For the dice to have landed in precisely this way when thrown would be a relatively rare event. Prior to my rolling the dice, the probability, given my evidence, that they would land in precisely this way would have been low. In spite of this, there's nothing particularly *surprising* about the way the dice landed – they had to land in some way and it might just as well have been this as any other. Suppose I shuffle a deck of cards and then examine the order in which the cards ended up. If the cards happen to be arranged in suits, say, then that could be regarded as surprising (though this could still be disputed). But, if there's no discernible pattern to the way the cards ended up, then surely nothing at all surprising will have happened. Nevertheless, for the cards to end up in the particular order that they did would be exceedingly rare – indeed, it is almost certain to be a unique event.

Lottery outcomes seem to provide another illustration of this general phenomenon. Suppose ticket #6255 is the winner. Suppose I hear the announcer read out this number and watch

a stranger go and collect the lottery prize. It's clear that I haven't witnessed anything surprising here. I knew from the start that some ticket was going to win, and it might just as well have been this ticket as any other. If ticket #6255 winning the lottery should come as no surprise to me then, given the initial premise, I cannot justifiably believe, just on the basis of the odds involved, that ticket #6255 has lost.

When it comes to my own ticket – ticket #5472 – things do, though, seem different. As I observed, it seems as though it really would be surprising to hear the announcer read out this number, and this may even partly explain the appeal of the first step of the lottery paradox. Even if I can't justifiably believe that ticket #6255 has lost, can I justifiably believe this about *my* ticket – ticket #5472? On reflection, I clearly don't have any special insight into the fate of my ticket that I lack when it comes to ticket #6255. In a pessimistic mood, it may seem as though the fact that I own a ticket gives me some special reason to think that it has lost but in reality that just isn't so. It's clear that I have no more reason for believing that ticket #5472 has lost than I do for believing that ticket #6255 has lost – if I can't justifiably believe the latter, then I can't justifiably believe the former either. If I can only justifiably believe a proposition when its falsity should be surprising for me, and all lottery beliefs have the same justificatory status, then it follows, as required, that no lottery beliefs are justified.

Similar considerations suggest that, while I might have a special interest or investment in certain lottery outcomes, the outcomes don't really differ as to how surprising they are – including the outcome in which ticket #5472 wins. After all, do we really want to insist that it should be *more* surprising for ticket #5472 to win than for ticket #6255 to win – even though the lottery is fair, the two tickets have exactly the same chance of winning etc.? The fact that ticket #5472 is my ticket might make its winning into a more *desirable* or *significant* or *lucky* event for me than ticket #6255 winning – but a more *surprising* event? If I shouldn't be surprised by ticket #5472 winning, then we have a far more direct argument for denying that I'm justified in believing that ticket #5472 has lost. If, on the other hand, we maintain that ticket #5472 winning really would be a surprising event, then we will have to make do with the slightly more circuitous argument outlined above.

None of the present considerations suggest that I can't justifiably invest a high *subjective probability* in the proposition that ticket #5472 has lost, and a low subjective probability in the proposition that ticket #5472 has won. Indeed, a suspension of belief, combined with a low subjective probability, seems to be the attitude that is perfectly tailored towards events that would be rare yet unsurprising. Prior to rolling the three dice, for instance, I shouldn't believe that they *won't* come up 11-7-4 – in one sense I should keep an open mind about how the dice are going to land. At the same time, if I was asked how likely it is that I'll roll 11-7-4, then I should regard that as highly unlikely.

IV Justification, Risk and Normic Support

Believing involves an element of risk. When I believe a proposition P, I am committed to P being true, and exposed to the risk that P is false. What it is for my belief to be justified, on one natural way of thinking, is for this risk of falsity to be properly managed or controlled. This simple idea provides some motivation for thinking that lottery beliefs *should* count as justified. After all, in believing that ticket #5472 has lost the lottery haven't I effectively managed the risk of falsity? Given my evidence, the belief is very unlikely to be false and very likely to be true. On further reflection, though, matters are not so clear-cut. While there is undoubtedly one sense in which my belief that ticket #5472 has lost runs a low risk of falsity, there is another sense in which it is at *considerable* risk of being false. While it is unlikely that the belief is false, very little would *need to happen* in order for the belief to be false. When a hand goes into the barrel to draw the winning ticket all that would need to happen, in order for my belief to be false, is for it to close around ticket #5472. Even if we suppose that a different ticket was in fact drawn, and my belief is true, it would have taken very little for the belief to have been false instead – if the barrel had been spun with slightly more vigour, or the person drawing the winner had rummaged a little longer or reached a little further, that may have been enough. Some sequence of perfectly mundane events is all that it would have taken for ticket #5472 to be the winning ticket.

For contrast, suppose again that I wander into a room for the first time, have a perceptual experience as of a red wall before me, and immediately come to believe that there is a red wall before me. Suppose the belief is true – there is a red wall before me and I am perceiving it in a perfectly straightforward way. In this case, it would have taken something *substantial* for my belief to have been false instead of true – not only would the wall need to have been a different colour, but my visual system would need to have been seriously malfunctioning, or there would need to have been deceptive red lighting installed in the room etc. In short, there would need to have been a substantial departure from how the world actually is. In the probabilistic sense, the belief that there is a red wall before me and the belief that ticket #5472 has lost may run a comparable risk of falsity – but there is another sense in which the latter belief is at much greater risk than the former.

Say that a body of evidence E is *safe evidence* for a proposition P just in case P could not easily have been false, given that E is true. Or, put in terms of possible worlds, say that E is safe evidence for P just in case, in all very close or similar possible worlds in which E is true, P is true. Say that a belief is safe just in case it is based on safe evidence. My belief that ticket #5472 has lost, based on the fact that there are 9999 losing tickets and only one winner, is not a safe belief. Whichever ticket happened to win the lottery, it could easily have lost – there are very similar worlds in which it does lose. Since any ticket could have won the lottery just as easily as any other, for any number n between 1 and 10000 there will be very similar worlds in which the lottery is run and ticket #n is the winning ticket. As such, there are very similar worlds in which my evidence holds and ticket #5472 is the winner. In contrast, my belief that there is a red wall before me, based upon my experience as of a red wall, is a safe belief. In all very similar worlds in which I'm having an experience as of a red wall before me, there really is a red wall before me. While there are possible worlds in which I am having this experience even though the wall is another colour – worlds in which

I'm hallucinating or colour blind or taken in by tricky red lighting – none of these lie close to actuality.

My belief that there is a red wall before me is safe, while my belief that ticket #5472 has lost the lottery is not safe. This may be one good sense in which the former belief is at lower risk of falsity than the latter (Pritchard, 2015, 2016). But this is not yet the end of the story. Suppose again that I wander into a room, have a perceptual experience as of a red wall before me, and immediately come to believe that there is a red wall before me – but suppose this time that the wall really *is* green and I really *am* undergoing a vivid colour hallucination. In this case my belief is clearly not safe. There is a very similar world – namely, the *actual* world – in which I'm having an experience as of a red wall before me even though there is no red wall before me. In spite of this, there is still a sense in which this belief is less risky than the belief that ticket #5472 has lost the lottery. While the belief does happen to be false, it *took* something substantial to make it so – it took a departure from how the world *would normally be*.

Suppose that possible worlds can be compared not just for their similarity, but for their *normalcy*. Say that a body of evidence *E* *normically supports* a proposition *P* just in case, in all of the most normal possible worlds in which *E* is true, *P* is true. Say that a belief is normically supported just in case it is based on evidence that normically supports it. My belief that ticket #5472 has lost is not normically supported by the fact that there are 9999 losing tickets and only one winner. In terms of normalcy, all of the possible lottery outcomes are on a par – no outcome could be considered more normal than any other. Amongst the most normal worlds in which the lottery is run there will be worlds in which ticket #*n* is the winning ticket for any *n* between 1 and 10000. Amongst the most normal worlds in which my evidence holds are worlds in which ticket #5472 is the winner. In contrast, my belief that there is a red wall before me is normically supported by the fact that I'm having an experience as of a red wall before me. In all of the most normal worlds in which I'm having an experience as of a red wall before me, there really is a red wall before me. While there are possible worlds in which the wall is another colour and I'm hallucinating or colour blind or taken in by tricky red lighting, none of these lie close to ideal normalcy.

My belief that there is a red wall before me is normically supported, while my belief that ticket #5472 has lost the lottery is not normically supported. This is another sense in which the former belief is at lower risk of falsity than the latter – and continues to be so even if it is, in actual fact, false. Even if the belief is actually false, it remains true that it *could not have been false* without the circumstances being highly abnormal. The belief that ticket #5472 has lost could be false without anything abnormal having happened.

I opened this section with the suggestion that securing justification for a belief is essentially a matter of managing or controlling the risk of falsity. While this idea might initially seem to support the claim that lottery beliefs are justified, on close inspection it does so only if we limit ourselves to a single *probabilistic* conception of risk. If we are open to a number of different conceptions of risk, then the justificatory status of lottery beliefs looks to be far less assured. By believing that ticket #5472 has lost the lottery, purely on the basis of the odds involved, there are a number of senses in which I have manifestly *failed* to manage the risk of falsity – managed the risk far worse than for many ordinary beliefs that I hold.

As I noted at the beginning of this section, by believing a proposition P I am committed to the truth of P and exposed to the risk that P is false. But to invest a particular *subjective probability* in P – even a subjective probability that is very high – is precisely to *avoid* committing to the truth of P, and avoid taking on this risk. If two subjects, one of whom believes P and one of whom assigns a high subjective probability to P, are confronted with the falsity of P, only the former needs to admit an error (Buchak, 2014, pp286-287, Fantl and McGrath, 2009, pp141). In taking it to be very likely, or 99.99% likely, that ticket #5472 has lost the lottery, I may be committed to the lottery having a particular kind of set-up – having a certain number of tickets, being fair etc. – but I am not committed to *anything* regarding the lottery's actual outcome⁵. As such, in order for this attitude to be justified, I do not need to manage the risk that ticket #5472 has won – my attitude does not expose me to any such risk.

V Justification and Epistemic Resilience

The final argument begins from the premise that a justified belief should exhibit some epistemic *resilience* or *inertia* in the face of contrary evidence. Suppose I believe a proposition P and then acquire some evidence against P – suppose, for instance, I read in a newspaper that P is false. This new evidence would usually have some effect on my epistemic position – but how my position ends up will depend upon the nature and the strength of whatever initial justification I had for believing P. If my initial justification for believing P was strong, the new evidence may still leave me with justification for believing P, albeit weakened. If my initial justification for believing P was moderate, then the new evidence may leave me in a balanced situation in which I ought to suspend judgment on P. If my initial justification for believing P was weak, then the new evidence might even leave me with justification for believing $\sim P$, albeit weaker than the justification that evidence of this sort would usually provide. But what if this new evidence left me with justification for believing $\sim P$ that was *every bit* as strong as the justification that such evidence would usually provide? This would tend to suggest that my initial belief in P had *no* justification whatsoever – that there was nothing for the new evidence to counterbalance. Put slightly differently, if my belief in P exhibits no epistemic resilience in the face of counterevidence, then this would tend to suggest that it is not a justified belief.

One possible way to gauge the *strength* of a person's justification for believing a proposition P is by counting up the number of independent, 'minimally reliable' information sources telling against P that the person would need to be exposed to in order for this justification to be neutralised (see the discussion of 'entrenchment' in Huber, 2009, section 4). One possible example of a

⁵ The question of exactly what a probabilistic judgment commits one to is delicate – and may not permit of a single, ubiquitous answer. Sometimes, to assign a probability to P is to estimate its current evidential probability – to make a judgment about how one's current evidence bears upon P, but nothing further. Other times, when P is conceptualised as one possible outcome that could be generated by a particular process, given certain specified initial conditions, then assigning a probability to P can implicate a judgment about the dispositions or propensities of that process. In this case, the assignment will be answerable to certain objective features of the world – though not, typically, to P itself. I won't pursue this further here (for some further discussion see Buchak, 2014, section 1).

minimally reliable source of information might be testimony from a stranger – a person about whom one has no prior information. Consider the following set-up (based on Huber, 2009, pp22). Suppose I believe P and then head out onto a crowded street asking passers-by whether P. Suppose that the first four people I encounter each tell me $\sim P$ and, at this point, I lose my justification for believing P, and am epistemically obliged to give up my belief and suspend judgment on the matter. We might say, in this case, that my initial justification for believing P had a strength of 4 – it took four pieces of contrary testimony to cancel it out. If this happened after only three people had testified to $\sim P$, then this would indicate that my initial justification had a strength of 3 and so on. Furthermore, if the first four people I encountered told me $\sim P$ and, at this point, I had justification for believing $\sim P$ instead of P, then this would indicate that my initial justification for believing P had been cancelled out *and then some* – it would indicate that my initial justification for believing P could not have had a strength any greater than 3. Finally, if I encounter just a *single* person who tells me $\sim P$ and, at this point, I have justification for believing $\sim P$ instead of P then this would indicate that my initial justification for believing P could not have had a strength any greater than 0 – which is just to say that I must have lacked any initial justification for believing P.

This picture is, of course, highly simplified and idealised – there are a number of potential complications that are being ignored. The epistemic force of testimony can, for instance, be affected by background evidence, as can the way in which testimony and other kinds of justification interact. I won't take it for granted that resilience in the face of contrary testimony provides a *perfect* guide to the justificatory strength of a belief – but I will assume that it can serve as a rough or approximate guide, provided interfering factors are not obviously present. In particular, I will assume that if a belief in P has no epistemic resilience – if a single piece of testimony telling against P is enough to furnish one with justification for believing $\sim P$ – then this is a good indication that the belief was not justified.

With this in mind, consider the lottery scenario once again. I know that there are 9999 losing tickets and only one winner. I know that the lottery has been drawn, but I've not yet heard anything about the result. In spite of this, I believe, purely on the basis of the odds involved, that ticket #5472 has lost. Suppose I then overhear two strangers talking and one says 'Did you hear the lottery result? Ticket #5472 was the winner'. Having heard this, I immediately acquire justification for believing that ticket #5472 has won. If I take the testimony at face value, and come to believe that ticket #5472 has won, then this would be a justified belief. By the above reasoning, this strongly suggests that I lacked any initial justification for believing that ticket #5472 has lost – that my justification for believing this had a strength of 0.

When I overhear this conversation, and acquire justification for believing that ticket #5472 won, it is worth noting that this justification is *every bit* as strong as my justification for believing anything else that I overhear. Suppose I eavesdrop on more of the conversation and I hear one of the strangers say 'This is the tenth straight year that the lottery has been run. They raised a record amount for charity'. I have just as much justification for believing that ticket #5472 won as I do for believing either of these claims. Imagine I do have some prior evidence that casts doubt on one of the claims – suppose, for instance, I dimly remember that the lottery was first held in the same year as my fifth wedding anniversary, and I know that this was only 8 years ago. In this case I would have somewhat less justification for believing that this is the tenth straight year the lottery has been run. But there is nothing that weakens my justification for believing that ticket #5472 has won, once I

acquire the testimonial evidence in its favour. The force of this evidence is not partly spent on having to overcome some pre-existing justification for believing that ticket #5472 has lost – there is nothing to overcome.

I have argued that my belief that ticket #5472 has lost the lottery, based purely on the odds against it winning, is not justified, as the belief has no epistemic resilience in the face of contrary evidence. No similar argument will show that I am not justified in assigning a high subjective probability to the proposition that ticket #5472 has lost. In general, a subjective probability assignment may be justified, even if it is highly epistemically sensitive to whatever further relevant evidence comes to light. Suppose I roll a regular six-sided die and it lands just out of view. At this point, I might assign a subjective probability of $1/6$ to the proposition that the die is showing 1 and, plausibly, I would be justified in doing so. Naturally, though, any evidence to the effect that the die is in fact showing 1, no matter how slight, would have a dramatic effect on the subjective probability that I ought to assign to this proposition. If someone tells me that the die landed 1, then the subjective probability I assign to this proposition should dramatically increase. This doesn't show, however, that there was anything wrong with my initial low assignment. What it illustrates is that the initial assignment was, in effect, a justified response to *ignorance* – to a *lack* of information.

What justifies me in assigning a subjective probability of $1/6$ to the proposition that the die is showing 1, is the fact that this represents 1 of 6 mutually exclusive and jointly exhaustive possibilities, and my evidence provides equal support to each. This assignment is nothing more than a stop-gap, quickly made obsolete by any evidence as to how the die landed. This is my epistemic position in the die case – but my position in the lottery case is, in all essential respects, exactly the same. After all, I have no more information about the lottery outcome than I do about the outcome of the die roll – it's just that, in the lottery case, there are *10000* possible outcomes over which to distribute my subjective probabilities⁶. It is my lack of information about the lottery outcome that justifies me in assigning a very low subjective probability to the proposition that ticket #5472 has won and a very high subjective probability to its negation – but this assignment is superseded once I acquire some evidence about the lottery outcome. This highlights another crucial difference between subjective probability and belief: A subjective probability assignment – even a very high one – can, in a manner of speaking, be justified by a *lack* of information. A belief, however, is never a justified response to a lack of information.

⁶ The fact that one's evidence is *neutral* with respect to a set of mutually exclusive and jointly exhaustive possibilities does not, in and of itself, oblige one to assign the same subjective probability to each – this is one lesson of the so-called 'problem of multiple partitions'. In the lottery case, given the information that the lottery is fair, it's plausible that one's evidence provides *equal support* for each of the possibilities – a stronger condition than mere neutrality, which can also be ensured by a breakdown in the comparability of evidential support. In cases in which each of a set of exclusive and exhaustive possibilities is equally supported by one's evidence, one is arguably rationally required to divide one's subjective probability evenly amongst them (for discussion see Smith, 2015).

VI Conclusion

In this paper, I've presented a series of arguments for denying that lottery beliefs are justified. If any one of these arguments proves effective, it follows that the lottery paradox provides no strong reason for rejecting the principle of closure. Part of my aim here has indeed been to safeguard closure against the threat that the lottery paradox seems to pose. But I've also tried to show that, when it comes to questions about the justificatory status of lottery beliefs, *it's not all about closure and the lottery paradox*. Even if we end up rejecting the closure principle for some other reason (the preface paradox perhaps) *still* we should deny that lottery beliefs are justified. The supposition that such beliefs are justified puts us at odds with plausible principles linking justification to knowledge, surprise, risk and epistemic resilience.

References

- Bird, A. (2007) 'Justified judging' *Philosophy and Phenomenological Research* v74(1), pp81-110
- Buchak, L. (2014) 'Belief, credence and norms' *Philosophical Studies* v169(2), pp285-311
- Derksen, A. (1978) 'The alleged lottery paradox resolved' *American Philosophical Quarterly* v15(1), pp67-74
- DeRose, K. (1996) 'Knowledge, assertion and lotteries' *Australasian Journal of Philosophy* v74(4), pp568-580
- Douven, I. (2002) 'A new solution to the paradoxes of rational acceptability' *British Journal for the Philosophy of Science* v53(3), pp391-410
- Douven, I. (2008) 'The lottery paradox and our epistemic goal' *Pacific Philosophical Quarterly* v89(2), pp204-225
- Dretske, F. (1971) 'Conclusive reasons' *Australasian Journal of Philosophy* v49(1), pp1-22
- Ebert, P., Smith, M. and Durbach, I. (2018) 'Lottery judgments: A philosophical and experimental study' *Philosophical Psychology* v31(1), pp110-138
- Fantl, J. and McGrath, M. (2009) *Knowledge in an Uncertain World* (Oxford: Oxford University Press)
- Foley, R. (1993) *Working Without a Net* (Oxford: Oxford University Press)
- Foley, R. (2009) 'Beliefs, degrees of belief and the Lockean thesis' in Huber, F. & Christoph Schmidt-Petri, C. eds. *Degrees of Belief* (Dordrecht: Springer)
- Friedman, O. and Turri, J. (2015) 'Is probabilistic evidence a source of knowledge?' *Cognitive Science* v39(5), 1062-1080
- Hambourger, R. (1980) 'Belief in miracles and Hume's Essay' *Noûs* v14(4), pp587-604

- Harman, G. (1968) 'Knowledge, inference and explanation' *American Philosophical Quarterly* v5(3), pp164-173
- Harman, G. (1986) *Change in View* (Cambridge, MA: MIT Press)
- Hawthorne, J. (2003) *Knowledge and Lotteries* (Oxford: Oxford University Press)
- Horwich, P. (1982) *Probability and Evidence* (Cambridge: Cambridge University Press)
- Huber, F. (2009) 'Belief and degrees of belief' in Huber, F. and Schmidt-Petri, C. eds. *Degrees of Belief* (Dordrecht: Springer)
- Ichikawa, J. (2014) 'Justification is potential knowledge' *Canadian Journal of Philosophy* v44(2), pp184-206
- Kroedel, T. (2012) 'The lottery paradox, epistemic justification and permissibility' *Analysis* v72(1), pp57-60
- Kyburg, H. (1961) *Probability and the Logic of Rational Belief* (Middleton: Wesleyan University Press)
- Kyburg, H. (1970) 'Conjunctivitis' in Swain, M. ed. *Induction, Acceptance and Rational Belief* (Dordrecht: Reidel)
- Lehrer, K. (1974) *Knowledge* (Oxford: Clarendon Press)
- Littlejohn, C. (2012) 'Lotteries, probabilities and permissions' *Logos and Episteme* v3(3), pp509-514
- Littlejohn, C. (2017) 'How and why knowledge is first' in Carter, A., Gordon, E. and Jarvis, B. eds. *Knowledge First: Approaches in Epistemology and Mind* (Oxford: Oxford University Press)
- Nelkin, D. (2000) 'The lottery paradox, knowledge and rationality' *Philosophical Review* v109(3), pp373-409
- Peacocke, C. (2004) *The Realm of Reason* (Oxford: Oxford University Press)
- Pritchard, D. (2007) 'Sensitivity, safety and anti-luck epistemology' in Greco, J. ed. *The Oxford Handbook of Scepticism* (Oxford: Oxford University Press)
- Pritchard, D. (2015) 'Risk' *Metaphilosophy* v46(3), pp436-461
- Pritchard, D. (2016) 'Epistemic risk' *Journal of Philosophy* v113(11), pp550-571
- Reynolds, S. (2013) 'Justification as the appearance of knowledge' *Philosophical Studies* v163(2), pp367-383
- Ryan, S. (1996) 'The epistemic virtues of consistency' *Synthese* v109(2), pp121-141
- Schick, F. (1966) 'Consistency' *Philosophical Review* v75(4), pp467-495
- Smith, M. (2010) 'What else justification could be' *Noûs* v44(1), pp10-31

- Smith, M. (2014) 'Knowledge, justification and normative coincidence' *Philosophy and Phenomenological Research* v89(2), pp273-295
- Smith, M. (2015) 'Evidential incomparability and the principle of indifference' *Erkenntnis* v80(3), pp605-616
- Smith, M. (2016) *Between Probability and Certainty: What Justifies Belief* (Oxford: Oxford University Press)
- Smith, M. (2017) 'Why throwing 92 heads in a row is not surprising' *Philosophers' Imprint* v17(21), pp1-8
- Sutton, J. (2007) *Without Justification* (Cambridge, MA: MIT Press)
- Turri, J. and Friedman, O. (2014) 'Winners and losers in the folk epistemology of lotteries' Beebe, J. ed *Advances in Experimental Epistemology* (London: Bloomsbury)
- Williamson, T. (forthcoming) 'Justifications, excuses and skeptical scenarios' in Dutant, J. and Dorsch, F. eds. *The New Evil Demon* (Oxford: Oxford University Press)