Rational Belief and Statistical Evidence: Blame, Bias, and the Law[1]

Dana Kay Nelkin

To appear in *The Lottery Paradox* edited by Igor Douven (Cambridge University Press)

1. Introduction: The Lottery Paradoxes and Statistical Evidence

It is rational to believe all sorts of things, even things about which we are not absolutely certain, and about which we admit that there is a chance that we are wrong. And we know that it is possible to have excellent evidence for propositions that are false. For example, I believe—rationally—that I will be in Munich next month, on the basis of my intention to be there, my having taken all requisite steps to enact my plan, such as making a flight reservation, and yet I admit that there is a small chance that I will not in fact be in Munich. Some small percentage of transatlantic flights are cancelled, there is a small chance of a natural disaster such as a volcano eruption that grounds all flights across Europe, and so on. And yet when it comes to some propositions, we can have apparently very strong evidence that makes their truth extremely likely—even more likely than that I will be in Munich next month—without its being rational to believe them. For example, there is a good argument that it is not rational for me to believe that my single ticket in a million-ticket lottery will lose, based on the evidence that 999,999 tickets will lose.[2] Is this correct? If so, how can we explain it?

One argument for this is based on the Rationality Version of the Lottery Paradox. In a nutshell: if it were rational to believe that my ticket will lose, it would, by parity of reasoning, be rational to believe that every other ticket in the lottery will lose, and yet it is also rational for me to believe that not every ticket in the lottery will lose. Thus, it would appear that it is rational for me to believe a set of things that I can see are inconsistent (for each ticket, that it will lose and that not every ticket will lose), and this is clearly a *reductio ad absurdum* of our assumptions. This set of claims together form what I call the Rationality Version of the Lottery Paradox.[3]

---

[1] An earlier version of this paper was presented at the Current Themes in Practical Philosophy Workshop at the Ludwig-Maximilians-Universität München, and I am very grateful for excellent feedback from the participants. Special thanks to Monika Betzler, Kathleen Connelly, Sam Rickless, and Anna Wehofsits for their very helpful comments on written versions of the paper. Finally, I am grateful to Tyler Burge for first bringing the Lottery Paradox to my attention as a graduate student working on free will and rational agency, and to Igor Douven for the opportunity to again bring together a variety of my interests in returning to it in the writing of this paper.

[2] See, for example, Harman (1968), Kaplan (1996), Nelkin (2000).

[3] More formally and completely:
(1*) It is rational for Jim to believe that his ticket (t1) will lose.
(2*) If it is rational for Jim to believe that t1 will lose, then it is rational for Jim to believe that t2 will lose, it is rational for Jim to believe that t3 will lose ... and it is rational for Jim to believe that t1,000,000 will lose.
(3*) It is rational for Jim to believe that t1 will lose ... and it is rational for Jim to believe that t1,000,000 will lose. (1*, 2*)
(4*) It is rational for Jim to believe that either t1 will not lose or t2 will not lose ... or t1,000,000 will not lose.
(5*) Propositions of the following form comprise an inconsistent set: (a) p1 ... (n) pn, (n+1) not p1 or ... not pn.
(6*) Jim recognizes that the following propositions comprise an inconsistent set: (i) t1 will lose . . . (n) t1,000,000 will lose, either t1 will not lose ... or t1,000,000 will not lose. (Adapted from Nelkin (2000))
(7*) It is rational for Jim to believe inconsistent things that he recognizes are inconsistent. (3*, 4*, 5*, 6*)

The absurd conclusion can be avoided by rejecting the initial claim that it is rational for me to believe that my ticket will lose. If we accept this solution to the paradox, then we must also accept that it can be rational for me to believe things such as that I will be in Munich next month that are less likely than things it is not rational for me to believe, such as that my ticket will lose in a million-ticket lottery. And the same goes for beliefs about the past, as well as about the future. For example, it is rational for me to believe that the vote originally scheduled for December on Theresa May's Brexit plan was postponed, even though newspapers sometimes make mistakes and I sometimes misread, for example, even if the chance of mistake is higher than that my lottery ticket will lose.

This solution to the paradox is controversial, though it appears to be gaining somewhat in popularity as of late.[4] To this point, philosophers have been much more likely to accept a parallel solution to the *Knowledge* Version of the Lottery Paradox. According to that version: if I know that my ticket will lose, it would, by parity of reasoning, be the case that I know that every other ticket in the lottery will lose, and yet I also know that not every ticket in the lottery will lose. Thus, it would appear that I can know a set of things that are inconsistent (for each ticket, that it will lose and that not every ticket will lose), and this is clearly a *reductio ad absurdum* of our assumptions.[5] A widely accepted solution to this Knowledge Version of the Paradox is to reject the assumption that I know that my ticket will lose. I simply do not *know* that my ticket will lose despite its likelihood of doing so, and this does not sound counterintuitive. But when it comes to the Rationality Version, people have been much more reluctant to reject the parallel assumption that it is *rational* for me to believe that my ticket will lose. They might say, "I admit that I don't *know* that my ticket will lose; but surely it's rational for me to believe it, given the high likelihood."

In reply, I have argued that there is a principled distinction to be made between my beliefs that I will be in Munich and that the Brexit vote was postponed on the one hand, and the ticket-buyer's belief that her lottery ticket will lose on the other hand. It is that the evidential support for the belief that her ticket will lose is exhausted by statistical evidence, whereas this is not the

---

(8*) It cannot be rational to believe inconsistent things that one recognizes are inconsistent.
(9*) (1*), (2*), (4*), (5*), (6*), or (8*) is false. (Adapted from Nelkin (2000))

[4] Admittedly, "gaining in popularity" should be understood relative to its starting point. Compare Redmayne (2008) to Buchak (2013), for example.

[5] More formally and completely:
(1) Jim knows that his ticket (t1) will lose.
(2) If Jim knows that his ticket (t1) will lose, then he knows that t2 will lose, he knows that t3 will lose ... and he knows that t1,000,000 will lose.
(3) Jim knows that t1 will lose ... and Jim knows that t1,000,000 will lose. (1, 2)
(4) Jim knows that either t1 will not lose or t2 will not lose ... or t1,000,000 will not lose.
(5) Propositions of the following form comprise an inconsistent set: (a) p1 ... (n) pn, (n+1) not p1 or ... not pn.
(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.

case when it comes to the other fallible beliefs, such as those about Munich and Brexit. My evidence in the Munich case is not merely statistical; it includes my awareness of my intentions and plans, my visual confirmation of a receipt for my plane ticket, and so on. My evidence in the Brexit vote postponement case includes my having read about it in several trusted online news sources and my having discussed it with others who have additional sources. Because this proposed solution to the Lottery Paradox draws a distinction between beliefs supported solely by statistical evidence and beliefs supported by non-statistical evidence, I call it the "Statistical Support Solution". It relies on this key Statistical Thesis (ST):

> (ST) For any proposition P, it is not rational to believe P solely on the basis of a high statistical probability of P.

In other words, it is not rational to form beliefs on the basis of what I will call "P-inferences", namely, inferences of the form "p has a statistical probability of n [where n is a very high number] → p."

Thus, (ST) nicely sorts the cases in question, while offering a solution to the Rationality Version of the Lottery Paradox. But is there independent reason for thinking that it is true? I believe that there is, and it ultimately rests on the nature of rationality. Even among those who offer very different solutions to the Lottery Paradox, many agree that rationality is a guide to the truth, and this is the starting point to an explanation for the fact that P-inferences are not rational ones. For believing solely on the basis of statistical evidence is inconsistent with one's being able to posit a causal or explanatory connection between one's belief and the fact that makes it true. In the lottery case, what explains the fact that it is not rational to believe that one's ticket will lose is that one cannot see a causal or explanatory connection between one's belief and the fact that makes it true. In fact, given the nature of the evidence, one can see that there is no such connection to be found. The role of rationality in one's belief-forming activities is to guide one's beliefs toward the truth, via reasons. There is a way in which this role is fulfilled in inferences to beliefs such as that I will arrive in Munich next month, but not in the case of P-inferences. In the everyday case like that of my belief that I will be in Munich or that there was a postponement of the vote on Theresa May's Brexit deal in the UK Parliament yesterday, the rationality of one's belief depends on one's being (rationally) committed to the proposition that there is an explanatory connection between one's belief and what the belief is about.

We can see the implications of this by supposing that I come to learn that my belief is false. Suppose that I find out that there was in fact a vote on the Brexit deal as originally scheduled after all. In that case, I would now reject something to which I was previously committed: that the newspaper reported accurately, or that there was a causal connection between the reporting and events in Parliament, or something else in my total evidential package. In recognizing the falsity of my belief, I realize that my evidence (although perhaps quite reliable) failed to connect to the truth in some way. Either the evidence was inaccurate, or it was not connected with the truth in the right way.

Importantly, I do not need to see myself as having been irrational or at fault in any way. But once I realize that my belief is false, I must reject part of what I took to give me reason, where this can include considerations in the foreground such as that I viewed the headline on my computer screen, or in the background, such as that I took it that the reporter witnessed the official act of postponement or had some other relevant explanatory connection to the event. In other words, in seeing that my reasons failed to guide me to the truth, I must reject something in my total "package" of reasons. Notice, in contrast, that this is not the case for P-inferences. If I believe that p (say, my ticket will lose) on the basis of a high statistical probability for p, and I find out that not-p (I won!), then there is nothing at all in my reasons to reject. I still believe the same odds were in effect, and I still believe that they made my losing extremely probable. I have no reason to think that my evidence failed to bear a connection to my conclusion that I previously thought it did. Learning that my belief is false puts no pressure on me to find some problem in my reasons. Thus, there is a way in which they are not "sensitive" to the truth, or at least to what I conceive of as the truth. Given the role of rationality as a guide toward truth, this lack of sensitivity to the truth in the case of P-inferences helps to explain why such inferences are not rational.[6]

So far, then, we have a solution that sorts cases in an intuitive way, and a rationale in terms of the connection between rationality and truth. Notice that this solution also helps explain subtle differences, as well as similarities, in the behavior that probabilistic beliefs and outright beliefs rationalize. On the one hand, the behavior that the merely probabilistic belief that one's ticket will win rationalizes overlaps substantially with the behavior that the outright belief does. If I believe my ticket will very probably lose, then I should not plan on quitting my job or altering my life in any substantial way. This is also true if I believe my ticket will lose. But there is a difference, as well. For example, if I believe my ticket will lose (as I would if I thought the stub would be torn up immediately), then I wouldn't be in a position to rationally buy it in the first place. But if I believe that it will very probably lose, it could still be rational for me to buy it. This point helps support the idea that lottery ticket buyers typically believe that their tickets will probably lose, but not that they do not typically believe that they will lose. For subtly different sets of behavior seem to be rationalized by the two beliefs. (In what follows, I will sometimes refer to beliefs that have probabilistic contents, such as that 'there is a 1/1000 chance that my ticket will lose' "probabilistic beliefs" and beliefs without probabilistic contents, such as that 'my ticket will lose', as "outright beliefs".)

Despite these advantages, the solution faces serious challenges. One source of resistance is precisely the kind of case with which I began: if it is rational to believe anything at all—such as that I will be in Munich—it is surely rational to believe a proposition that appears to have an even higher likelihood of being true, such as that my ticket will lose in a million-ticket lottery. Further, if rationality is ultimately aimed at truth or accuracy, then this solution seems problematic insofar

[6] It might be thought that it is the fact that the lottery belief in question is about the future that undermines the rationality of a P-inference in this case. But we can imagine that the lottery has already been conducted and the result kept secret for a time, and in that case, as well, the same explanation is effective. Similarly, as we will see, statistical evidence about past crimes poses a particular challenge in trial law.

as it denies rationality to a belief that I can see is *more* likely to be true than one that is rational.[7] This is to understand the connection between rationality and truth in a way that is different from the one I set out earlier.

One way of extending this idea is to reject talk of "outright" belief, or "binary" belief (i.e., you believe it or you don't) altogether as capturing anything epistemically important.[8] While speaking of belief might be useful, it is an inherently imprecise way of talking, and what really matters is "graded beliefs" or "credences", that is, degrees of confidence in propositions. And these can be as precise as you like. I can have a .0001 credence that my lottery ticket will win, say. And once we replace talk of rational belief with talk of rational credence, we can solve the Lottery Paradox by offering epistemic principles that govern the coherence of sets of credences, rather than requiring simple logical consistency as we do when it comes to belief.[9] By making this move, proponents avoid any solution to the Lottery Paradox that presupposes rational (binary) belief, including one that appeals to (ST).

While I briefly address these points at the end of the paper, my main aim here is to build a positive case for (ST) by showing that it has great explanatory power in a number of domains, including the conditions for appropriate blame, stereotype-based beliefs, and standards of proof in the law. In each domain, I argue that (ST) can either help solve a puzzle or best explain an otherwise unexplained phenomenon. While others have noted connections between (ST) and each of the three domains, my aim here is to bring these together to show the explanatory power of (ST), and to show how the thesis fits with and mutually supports the most satisfying accounts in each domain. (ST) has competition in each of the three domains, and although the comparative assessments are of necessity incomplete, I set out reasons for nevertheless preferring its inclusion in each area.

2. Bias and Belief

Some beliefs involving racial or sex-related stereotypes strike us as deeply problematic. And yet, as has been pointed out, some stereotypes are rooted in empirical facts. It will help to have some examples to work with. Let us begin with what I will call *The Cosmos Club Case*. John Hope Franklin, renowned historian, professor, and best-selling author, describes an incident that occurred in the Cosmos Club in Washington, D.C. on the eve of his receiving the Presidential Medal of Freedom from President Clinton. As a member of the club, he gave a small dinner party for friends. He writes,

[7] See, e.g., Douven (2008).

[8] E.g., See Christensen (2004) and (xxxx).

[9] While some (e.g., Ho Hock Lai (2008)) distinguish between the belief that P is .n likely to be true and the .n-credence that P, others take these to come to the same thing. For example, Moss (2018a) writes that "to believe that it is at least .3 likely that Jones is shoplifting is to have at least .3 credence that Jones is shoplifting" (183). Moss's illustrative example here seems particularly problematic in that the "at least" appears in the propositional content of the belief, but not in the content of the credence.

> It was during our stroll through the club that a white woman called me out, presented me with her coat check, and ordered me to bring her coat. I patiently explained that if she would present her coat to a uniformed attendant, "and all of the club's attendants were in uniform," perhaps she could get her coat. (2005, 340)

This is obviously very disturbing, and something has gone very wrong here. At the same time, as Tamar Gendler points out, the woman was not obviously guilty of making a simple base-rate error. According to Gendler,

> Franklin had been the Cosmos Club's first black member, and was still one of its very few. By contrast, nearly all of the club's numerous attendants were men of African descent. So when the woman was presented with the visual experience of a black man in the club's lobby, she endorsed an empirically well-supported hypothesis—one that took full account of prior probabilities. The likelihood that a black man present in the Cosmos Club was a member of the staff rather than a member of the club was very high—high enough, perhaps, to make it rational to assume that even though he was wearing a suit rather than a uniform, he was nonetheless an employee rather than a host. (2011, 35)

Now there are many things going on in this case, and arguably not wearing an attendant's uniform should provide an additional relevant base rate category for guests of the club. But setting that point to the side, we can certainly imagine similar sorts of situations in which being a member of a certain category (whether one of apparent race or sex, say) makes one considerably more likely to possess some other significant feature (such as occupation or wealth or particular kind of criminal record).

There are at least two ways to explain what is problematic here. One is to identify the mistake as related to the woman's belief, such as that she was not in fact rational in assuming that he was an employee rather than a host, and the other is to say that though the belief on which she acted was itself unproblematic, it was wrong in some other way for her to act as she did.

Cases like this one and others have led many to think that even if she also made mistakes in acting the way she did, her *belief* about Franklin was problematic in itself. And this seems right. (ST) is well-poised to explain what is problematic about the woman's belief.[10] She did not see Franklin wearing a uniform that would identify him as an employee, she did not witness him taking or returning others' coats, and she did not recognize any other reason specific to him to suggest that he was working at the club. Rather, her belief seems to have been based entirely on a statistical generalization related to his appearance as belonging to a racial category. (ST) denies that this belief is rational. The woman might rationally believe that it is *likely* that Franklin is an employee. But that probabilistic belief licenses behavior that is different from the behavior

[10] Note that Georgi Gardiner (forthcoming) sees this possibility. Notably, she seems to simply assume (ST) in doing so. She does not rely on this point, however, to explain what she takes to be problematic about biased beliefs of this kind.

licensed by the belief that he *is* an employee. To be sure, just as in the lottery case, there is overlap in the behavior that the two beliefs (one probabilistic and one not) license. But there are also differences. And it seems clear that if she had been epistemically careful and formed only a probabilistic belief, and had acted rationally on that belief instead, she would not have given him her coat check. She might instead have sought out further evidence, or made her way to the coat check counter.

Thus, (ST) figures neatly into a plausible explanation of what has gone wrong in the Cosmos Club Case. (ST) need not explain all that has gone wrong. For example, neither the probabilistic belief nor the non-probabilistic belief regarding Franklin's employment would justify "ordering" him to get her coat. But I think that even in a hypothetical variant of the case in which the woman treats people she views as employees in a more appropriate manner, we would still find it problematic for her to assume that Franklin is an employee. And (ST) helps explain what has gone wrong here.[11]

So far, so good for (ST), but, alas, I cannot leave things here. For an alternative approach to this sort of case has attracted a great deal of support in recent years. Like (ST), it denies that it is rational for the woman to believe that Franklin is an employee of the club. But rather than appeal to the idea that she mistakenly bases her belief entirely on statistical evidence, it appeals to the idea of moral encroachment. This idea, as applied to the case at hand, is that it was not in fact rational for the woman to believe as she did because what it is rational for one to believe—or what one is justified in believing—depends on the *moral stakes* in the situation. And because so much was at stake in this context, the standard of justification is much higher than it would have been in a low stakes case.[12]

Now the idea that there is such a thing as moral encroachment certainly has its critics. Here I will simply try to show that such a radical rejection of a foundational pillar in epistemology, namely, that epistemic norms and standards are sensitive to epistemic considerations and not practical or moral ones, is not required to explain at least some of the phenomena that have motivated it. As we've seen, when it comes to cases like the Cosmos Club, (ST) together with principles bridging probabilistic belief and action, can explain what is problematic about the woman's beliefs in the case.

But advocates of the moral encroachment explanation point to related cases that moral encroachment appears to be able to explain while (ST) simply has no application. In particular, as Moss (2018) and Gardiner (forthcoming) have pointed out, even some probabilistic beliefs seem problematic when their content is about race. For example, consider these cases.

[11] In addition, sometimes moral norms themselves make ineliminable reference to epistemic ones. So, for example, it might be that there is a moral norm to use substantial care and attentional resources when in situations in which risk of harm is salient, to avoid making epistemic mistakes. Thus, it might be that in this case, there is *both* an epistemic and a moral failure. See Nelkin (in preparation) for discussion of the relation of moral and epistemic norms and responsibility.

[12] See Kim (2017) for a recent survey on the more general phenomenon of pragmatic encroachment.

In the case I will call "Tipping," a server at a restaurant decides on the basis of the skin color of a diner that he will not tip well. And in another racial profiling case I will call "Sidewalk Switcher," a pedestrian sees someone she takes to be of a different racial group than herself walking toward her on the sidewalk and crosses the street in order to create more distance from him.[13]

If moral encroachment, or some other explanation, could account for what is problematic in all three cases while (ST) could only explain one, then that would be an advantage over the explanation that appeals to (ST) in the Cosmos Club Case. And this would in turn undermine the argument for (ST) based on its explanatory role in the Cosmos Club Case. Now I note that this point would not necessarily be decisive, since unity of explanation is one explanatory virtue among others. Perhaps (ST) remains a good explanation in one kind of case, even if other explanations are needed in other cases. Nevertheless, the comparison to the moral encroachment explanation bears further scrutiny.

The first thing to note is that at least sometimes cases like Tipping involve outright belief. I am not as confident as Moss, for example, about people's scrupulousness in sticking to their probabilistic beliefs in such cases. In fact, the existence of cases like the Cosmos Club Case suggest that it happens on occasion that people move from general probabilistic evidence to outright belief about individuals. And in at least some variants of Tipping, it seems that it would not make sense to treat a person of color as a bad tipper if one were careful to stick to probabilistic beliefs. Even if we were to confine the scope of reasons to the instrumental, the relatively low cost of providing good service with a, say, 90% chance that one will still receive a bad tip is not obviously sufficient to support one's providing bad service. So bad service might be rationalized and explained by the (mistaken) outright belief that the tip will be bad.

It is easier to make a strong case that the belief in Sidewalk Switcher is a genuinely probabilistic one. People play the odds in such cases. As Moss points out, when one switches sides of the street to avoid a pit bull rather than a person, it makes sense to describe the case as one in which one is playing the odds. Though I will grant that some cases are like this, I think the pit bull case is an interesting comparison that shows that even in Sidewalk Switcher, there might be more than just a probabilistic belief. In the pit bull case, one might think there is only a statistical probability that the pit bull will attack *on a given occasion*. But it might be because one has attributed—outright—a disposition to attack, that has a chance of being masked on a given occasion. If something similar were going on in the Sidewalk Switcher case, then though there would be a probabilistic belief about the odds of being assaulted, there would be an unsupported outright belief about a disposition.[14] But let us set aside this variation of the case, and grant that in some cases people merely have probabilistic beliefs.

The option I favor here is to acknowledge that in cases of *pure* probabilistic belief, where agents are careful not to adopt unjustified outright beliefs about individuals, there is simply not

[13] Both of these are adapted from cases described in Moss (2018a).
[14] See Sally Haslanger (2011), p. 200.

a problem with the belief itself. But first note that it is not easy to come up with a clear case of this kind. And, further, even if we do have a clear case, we need not accept that it is permissible to cross the street on such grounds. There are two reasons for this. First, the precise content of the probabilistic beliefs might be wildly unjustified epistemically, and there is good reason to suppose that this is often the case.[15] Second, even rational (and true) beliefs do not *individually* justify behavior. Whether any given behavior is justified depends on what else one justifiably believes in a given situation, the risks one's actions pose to others, and so on.

Thus, I take it that there is a unifying explanatory package that accommodates all three cases, as well as noted variants of each. It is one that appeals to (ST) directly to explain some cases, together with the idea that probabilistic and outright belief can rationalize different behavior. And it appeals to a more general recognition that beliefs rationalize or justify behavior holistically and not individualistically. Thus, while moral stakes do not affect epistemic norms, they *do* affect what beliefs are required for what sorts of actions, and can even affect the *level* of epistemic justification that is required before acting. In other words, they do not enter into what makes a given belief justified or rational, or what makes a belief justified to a particular degree. But they do enter into what makes a particular piece of *behavior* justified when one possesses particular rational beliefs.

Interestingly, this is a point well-noted in the context of blame and punishment (phenomena to which we will return shortly). For example, defenders of free will skepticism sometimes note that the arguments against skepticism must clear an especially high bar because there is so much at stake if we are wrong.[16] They do not mean that free will defenders' arguments are *worse* than they otherwise would be because of the stakes or that they provide less in the way of justification; rather, they mean that they need to add even more and better arguments in order to justify the kind of *action* they are claimed to justify.

Thus, once we adopt this plausible holistic picture of the justification of behavior, the thesis of moral encroachment becomes unnecessary. As for (ST), we can then see the thesis as one important component of a package of principles that together explain a wide variety of phenomena. Probabilistic beliefs, together with other beliefs and morally salient situational features license some, but not all, of the same behavior than do non-probabilistic beliefs.[17]

---

[15] See, for example, a classic set of general findings in Kahneman and Tversky (1996), and a more specific set of findings in the estimation of probabilities of criminal victimization in Quillian and Pager (2010).

[16] See Pereboom (forthcoming), for example.

[17] It is also worth noting two recent alternative suggestions to moral encroachment. Gardiner (forthcoming) suggests that at least one comprehensive explanation for the various kinds of cases appeals to the kind of understanding within which the apparently problematic beliefs are embedded. She argues that when they are embedded in a racist understanding, say, that is what makes them problematic. While I wholeheartedly agree that the racist understanding is morally problematic, this explanation does not depend on the purely probabilistic beliefs themselves being morally problematic. Insofar as this might be seen as a way of explaining *away* our intuition that there is something problematic about the purely probabilistic belief, it is plausible, however. A second alternative suggestion is made by Moss (2018a). She argues that even the credences associated with probabilistic evidence are epistemically deficient in that they do not constitute knowledge, *in this sort of case*. And this, in turn, is because the epistemic agents in the relevant cases cannot rule out relevant alternatives. There is

3. The Proof Paradox and the Law

Turn now to a much-discussed puzzle in trial law concerning standards of proof, sometimes known as the Proof Paradox. Consider the following legal cases, each adapted from real cases.[18] Judith Jarvis Thomson describes the following case, that I will call Bullets, as follows:

> Two people, Tice and Simonson, both hated Summers and wished him dead. Summers went hunting one day. Tice followed with a shotgun loaded with ninety-five pellets. Quite independently, Simonson also followed, but *he* had loaded his shotgun with only five pellets, that being all he had on hand. Both caught sight of Summers at the same time, and both shot all their pellets at him. Independently: I stress that there was no plot or plan. Only one pellet hit Summers, but this one was enough: it hit Summers in the head and caused his death. While it was possible to tell that the pellet which caused Summers' death came either from Tice's gun or from Simonson's gun, it was not possible to tell which (1986, 200-201).[19]

This case presents us with a puzzle, because it is highly likely that Tice was the actual murderer, and, yet, most people feel quite uncomfortable at the idea of convicting him on this basis. Courts, too, often disapprove of statistical evidence of this kind.[20] But now compare the case to one in which we have eye-witness testimony, identifying Tice. Even if we learn that eye-witness testimony is wrong in a similar percentage of cases, this does not engender the same reluctance to convict.[21] Again, then, we face a puzzle, and notably, it features a case in which we have only statistical evidence, compared to a case in which we have evidence that is not solely statistical.

Now one might think that the reason statistical evidence here is not sufficient for convicting Tice is simply because the standard of proof is too high: in a criminal case of murder, the standard is "proof beyond a reasonable doubt." There is, as yet, no puzzle here. But we can imagine a variant of the case in which the number of bullets Tice shoots is even higher—as high

much to say about this suggestion as there are a number of controversial commitments to this framework. Here I note that it is not clear why we should see lack of knowledge as so important, as opposed to rational belief. Moss suggests that the main upshot of attributing a lack of knowledge is a lack of assertability. But though I agree with Moss that "it sounds bad" for the pedestrian in Sidewalk Switcher to assert "You are more likely to steal my purse than anyone across the street," the badness of saying this to someone can easily be explained without appealing to a lack of knowledge. There are many reasons it would not only sound bad but *be* bad to assert this, including what it might convey in the circumstances, not to mention that we do not usually track down strangers to inform them of the statistical likelihoods of their having certain features given their group membership.

[18] The first is from Thomson (1986). See also Cohen (1977).

[19] If, like me, you don't think that there is a difference in the *degree* of culpability between Tice and Simonson because they both acted on identical intentions to kill, note that there are variants of the case that have it that there is only statistical evidence that a particular defendant even formed the intention to commit a crime. The important thing for our purposes is that the only evidence is statistical.

[20] See Ho (2008) for an overview of judicial decisions.

[21] This is not to say that we ought to simply accept eye-witness testimony.

as you like. And there still seems something problematic about convicting him. But we need not even turn to an unlikely scenario, for, as many have noted, the puzzle arises even in tort law, where we find a lower standard of proof, namely, that of "preponderance of evidence". So, consider this case, also described by Thomson, that I will call Red Cab:

> Mrs. Smith was driving home late one night. A taxi came towards her, weaving wildly from side to side across the road. She had to swerve to avoid it; her swerve took her into a parked car; in the crash, she suffered two broken legs. Mrs. Smith therefore sued Red Cab Company. Her evidence is as follows: she could see that it was a cab which caused her accident…and there are only two cab companies in town, Red Cab (all of whose cabs are red) and Green Cab (all of whose cabs are green), and of the cabs in town that night, six out of ten were operated by Red Cab….although she could see that it was a cab…she could not see its color, and there were no other witnesses to the accident—other than the driver himself, of course, but he has not come forward to confess.[22]

This time, we have statistical evidence that makes it more likely than not that the cab in question was red, and that would seem to meet the preponderance of evidence standard of proof. And yet, here, too, we feel discomfort at the idea of assigning liability to Red Cab Company.

Finally, consider a variant of this sort of case, called "Gatecrashers", as described by David Enoch et al:

> [I]t is uncontested that of, say, a thousand people attending a stadium event, only ten purchased tickets. If an individual—call him John—is sued, or, even more clearly, if he is prosecuted, then finding against John merely on the strength of the (very strong!) statistical evidence here seems to be inappropriate, even when convicting on the strength of a probabilistically equivalent piece of direct individual evidence (say, a videotape) seems perfectly fine (2012, 207).[23]

(ST) provides an elegant solution here. It is clear that our discomfort in convicting Tice or John or reaching a verdict of liability in Red Cab is driven by the fact that we have only statistical evidence in each case. According to (ST), if this is the only available evidence, then it is not rational to believe that Tice or John is guilty or that Red Cab is liable. Putting this implication of (ST) together with a requirement that legal fact-finders—jury members or judges—actually believe, rationally, on good grounds, in the guilt of defendants before reaching guilty verdicts, we have a powerful explanation of that discomfort. And there is good reason to think that we do require and value actual belief in guilt, and not just in the proposition that the defendant is probably guilty. Ho Hock Lai helpfully sets out a number of articulations of this standard by the courts. For example, the Supreme Court of Massachusetts wrote in *Sargent v Massachusetts Accident Company* that to make a positive finding there must be "actual belief in its truth", and Dixon J held in the Australian High Court case of *Briginshaw v Briginshaw* that "when the law

[22] Compare *Smith v. Rapid Transit, Inc.* and the hypothetical variant sometimes called the Blue Bus case.

[23] See David Kay (1978).

requires the proof of any fact, the tribunal must feel an actual persuasion of its occurrence or existence before it can be found."[24] (Whether we *should* require actual belief is an interesting question, but it appears that we do, and I think that there are good reasons for it in the end.) In sum, given that we have only statistical evidence in these cases, (ST) entails that it is not rational to believe that the parties in question are guilty or liable, respectively. And given that we require rational belief for finding of guilt or liability, it follows that the statistical evidence alone does not in fact suffice for legitimate conviction.[25]

But again, we cannot leave things here. For, in this domain, there are also competing explanations that likewise offer to validate our resistance to the sole use of statistical evidence. And, in addition, there are arguments that purport to show that our resistance is misplaced in the first place. While I do not here have space to do a comprehensive survey, it is worth examining an important example of each alternative.

Turn first to Thomson's own suggestion, which I believe has at times been mistakenly conflated, or lumped together, with the one just given that appeals to (ST).[26] So it is especially important to examine and assess. Importantly, Thomson quite explicitly rejects (ST), writing, for example, that "Alfred has *good reason* (emphasis mine) to think Bert will lose the lottery he entered, for Alfred is entitled to belief it .95 probable that Bert will lose the lottery" (1986, 210). Here and elsewhere Thomson makes it quite clear that outright believing that one's ticket will lose is *rational*; but the belief has a different and *additional* feature that makes it unsuitable for undergirding a verdict in the cases, namely, that it would be a kind of "luck" if it turned out to be true (214).

How are we supposed to understand "luck", according to Thomson? "Luck" is understood as in opposition to the idea that the evidence in question bears a causal relationship to the belief. And, in turn, to the extent that the fact in question is causally necessary for the existence of the evidence, the evidence "guarantees" the truth of the belief. For example, Thomson imagines the following variant of the lottery situation. Bertha bought a ticket in the lottery. But Alice witnesses the ticket seller tearing up Bertha's ticket stub right after Bertha bought the ticket. In this case, Alice believes that Bertha's ticket will lose, and not only does she have good reason, it is also not a matter of luck that her belief is true relative to her evidence. For, as Thomson explains, "…the fact which was Alice's reason for believing that Bertha would lose the lottery caused, and thereby guaranteed, that Bertha would lose it" (208).

Thomson takes it that an "anti-luck" requirement must be fulfilled if one's reasons or evidence is good *enough* to fulfill the so-called "third condition" (after truth and belief) on *knowledge*. Thomson makes this even stronger claim: "What marks A as having good enough reason for believing that p is true [to count as knowledge]? I suggest it is enough if A takes himself to have a guarantee that p is true, and if, also, what he takes to guarantee this does guarantee

[24] See Ho (2008), p. 106-107.
[25] See also Ho (2008), chapter 3 ("Epistemology of Legal Fact-Finding"), for a detailed argument in the same spirit.
[26] See, e.g., Wheeler (2007), Buchak (2013).

it" (212). Matters are subtle here. For Thomson clearly thinks that we can have rational beliefs that are not based on evidence that is perceived as guaranteeing truth. But the special kind of reasons we need to have to have *knowledge* requires that we see them as truth-guaranteeing. Thus, insofar as she is addressing the Lottery Paradox, she is at most addressing the Knowledge Version.

We might ask at this point how this is all related to the Proof Paradox. And it turns out that what most interests Thomson for this purpose is not knowledge itself, but rather the conditions for *assertion* of knowledge. For the delivery of a verdict is not identical to, but appears to be closely related, in Thomson's view, to the assertion of knowledge of guilt. When jurors find a defendant guilty, they must take themselves to have not *only* good reason to believe the defendant guilty, but also "reason of a kind which would make it not be just luck for the jury if its verdict is true." Their verdict would be unjust otherwise (214).

As Thomson sums up her solution:

> What would make it not be just luck for the jury if what it declares true is true? A guarantee. I suggested that individualized evidence for a defendant's guilt is evidence which is in an appropriate way causally connected with the (putative) fact that the defendant is guilty, and hence (putatively) guarantees the defendant's guilt; so to require individualized evidence of guilt just is to be requiring a guarantee. (214)

Thus, Thomson's approach does not appeal to (ST). Rather, it appeals to a distinct principle based on a kind of anti-luck condition for knowledge, which we might call (LKT) for the Luck Knowledge Thesis:

> (LKT) In order to have a reason *of the right kind* for one's non-probabilistic belief that p to constitute knowledge, one must not believe only on statistical evidence.

One reason that (ST) and (LKT) have been lumped together is that they share some key features, such as entailing that pure statistical evidence for a non-probabilistic belief falls short of some ideal. But they are also quite different, including in what the ideal in question is. Most notably, (ST) is a thesis about rational belief; (LKT) about knowledge. Further, the rationales that have been given for each are importantly different. Where both offer the idea that one must be able to posit a certain connection between one's evidence and belief, Thomson requires a causal and truth-guaranteeing connection; the requirement behind (ST) is that one be able to posit an explanatory connection, which might be causal, and needn't be guaranteeing.

In comparing the two explanations, it is worth noting that focusing on rational belief instead of knowledge has at least two advantages. The first is that (ST), unlike (LKT), is able to provide a solution to the Rationality Version of the Lottery Paradox, and is able to help explain other phenomena, such as the Cosmos Club and other bias cases in which the standard of knowledge does not appear to arise. Thus, (ST) has wider explanatory scope.

Second, (LKT), with its emphasis on knowledge, faces a special challenge in explaining how even individualized evidence that intuitively falls short of knowledge could be insufficient for liability in Red Cab. For it is implausible that knowledge, or even the assertion of knowledge, is associated with a preponderance of evidence standard in the first place.[27] Thus, (ST), which only requires that a minimal standard of rational belief would be required in both the Red Cab and Bullets cases in order to provide an explanation of what has gone wrong with mere statistical evidence again seems better placed to explain both cases in a unified way.

Finally, it is important to note that even if we agree that knowledge is the important epistemic state in play, the idea that it requires a guarantee seems implausibly strong. Perhaps this language could be relaxed so as to bring the rationale for (LKT) closer to that of (ST), by requiring only that one have reasons that one can posit have an explanatory connection to the truth of one's belief. But interestingly, this might still be more contentious when it comes to knowledge than when it comes to rational belief.

In sum, while there is something similar in spirit between (ST) and (LKT), namely, the idea of positing a causal or explanatory connection between reasons and belief, I hope to have shown that the most plausible application of this idea when it comes to both epistemology and legal verdicts is best captured by (ST) and its rationale, which requires for rational belief the ability to posit an explanatory connection between one's reasons and one's belief.[28,29]

I now want to turn briefly to a wholly different approach, on which our discomfort with the use of only statistical evidence in reaching verdicts, while perhaps understandable, is ultimately misplaced. David Enoch et al (2012) ask a reasonable question: even if there is an epistemological distinction to be made between beliefs based on the two kinds of evidence, why

[27] Thomson addresses the issue, but only by moving entirely away from an appeal to knowledge, and offering a standard that cannot be understood without an unintuitively high level of sophistication: "the friend of individualized evidence may be taken to say that the jury must believe it is more probable than not that the defendant is guilty *because* of believing it is more probable than not that there are facts available to it which guarantee that the defendant is guilty" (215). The standard is also an odd one, in that if there are facts available which guarantee guilt, then it would seem that the highest standard would thereby be met already.

[28] There are a number of other attempts to solve the proof paradox by appeal to the nature of knowledge. See Redmayne (2008) for an excellent survey up to that point, Enoch et al (2012) and (2015) for discussion of conditions of knowledge in relation to the paradox, and see Moss (2018b), among others, for a recent appeal to knowledge that addresses both standards of proof. (See note xx.)

[29] It is worth noting that (LKT) and (ST) are species of a single genus, namely, epistemic explanations. Both appeal to an epistemic deficiency of some sort (whether [with respect?] to rational belief or [with respect to?] knowledge) to explain what is problematic about relying solely on statistical evidence. There are other explanations of different kinds, however. For example, there are consequentialist explanations, such as that allowing prosecutors to rest on statistical evidence will have a dampening effect on their pursuit of additional evidence that might end up being either exonerating or condemning. (See Ho (2008) for examples.) Though I don't have the space to do justice to all of these, I think that, even if I imagine that in allowing evidence to rest wholly on statistical evidence we do not encounter any of the identified bad consequences, I am still uncomfortable with the idea that a verdict would be based solely on such evidence. I realize that a more complete defense of (ST) would involve a more extensive survey of alternatives than I can offer here. But I hope that I've showed the explanatory appeal of (ST), and its comparative advantage to at least one other species of explanation.

should this have any value in the law?[30] To make this vivid, they ask a further question: if we were choosing a system in which our children would live, would we choose one that, say, required knowledge (or some essential ingredient in knowledge) or one that was overall more accurate? The answer that we ought to choose accuracy. Their suggestion is not that accuracy trumps everything, but that it does trump epistemic value. If this is right, then even if (ST) is true, it would have no implications for legal practice.

My approach to this challenge has two parts. The first is to point out that, strictly speaking, the dilemma they set is between what they take to be a necessary condition of *knowledge*, namely, counterfactual dependence between belief and proposition believed, and accuracy. I take the point that a necessary condition of knowledge simply might not seem to be the most important value even among epistemic values. Yet an exclusive choice between rational belief on each occasion on the one hand and accuracy in the aggregate on the other hand is quite a different choice, however, and while a full adjudication is beyond the scope of this paper, I note simply that it is less tempting on its face to choose accuracy.[31] The second point is more targeted to present purposes. Even if it were to turn out that we should not distinguish between statistical and non-statistical evidence for purposes of legal verdicts, (ST), together with its rationale, would still serve as an excellent explanation of our current practices. The aim of achieving "actual persuasion" or "actual belief in its truth" when it comes to verdicts is clearly an appealing one—to the judges cited earlier and, I believe, to citizens who reflect on the matter. Insofar as we adopt this end, (ST) offers a natural explanation of our current practice and attitudes.[32]

4. Blame

Turn finally to a different, but related, phenomenon, that of blame. In criminal cases, it seems that not only are we reluctant to base guilty verdicts solely on statistical evidence, but we are reluctant to blame defendants when that is our only available evidence. And in ordinary cases, the same seems to be true. Consider this case offered by Lara Buchak (2013):

> You leave the seminar room to get a drink, and you come back to find that your iPhone has been stolen. There were only two people in the room, Jake and Barbara. You have no

[30] It is worth noting that Enoch et al ultimately do identify an "incentive-based" reason for making a distinction between the two kinds of evidence, while claiming that the epistemic distinction is explanatorily inert in the legal context. Others have questioned whether rational belief, or even whether belief is essential to verdicts. Se Ho (2008) for a helpful overview of this debate.

[31] Among other things, a full adjudication would need to assess the nature of each trade-off. In the case of a choice between a necessary condition of knowledge and accuracy, we need to look at the details of the condition of knowledge and exactly how to understand accuracy. For example, on Enoch et al's condition of knowledge, that condition arguably fares at least as well as accuracy when it comes to false positives (though not necessarily when it comes to false negatives). Because Moss embraces "probabilistic knowledge", the trade-off will look quite different.

[32] Though a full treatment would require an article of its own, I believe that this second strategy can be used as a partial answer to the use of DNA evidence, which itself turns out to be probabilistic in ways that aren't universally known.

evidence about who stole the phone, and you don't know either party well, but you know (let's say) that men are 10 times more likely to steal iPhones than women. (294)[33]

Would you blame Jake without collecting more information? I take it that we would withhold blame if our only evidence were statistical in this way, and that this would be the reasonable thing to do. Once again, (ST) seems to provide just the right explanation here, again, together with a bridge premise: Blame entails the belief that the person performed the wrong act (or omitted to fulfil an obligation), and that the person culpably did so. This explains our reluctance to say that blame would be appropriate under the circumstances. In fact, on what I take to be the most plausible conception of blame, the belief that the target of blame has engaged in culpable wrong-doing is part of what it is blame.[34] Accounts of blame that avoid attributing a belief in wrongdoing and blameworthiness to blamers face the serious challenge of over-inclusivity. For without the constraint of this belief, other accounts, such as those that appeal only to the emotion of anger, would seem to count as blame cases that are intuitively not blame. One might be angry, for example, without *blaming*, as one might be angry that a person behaved in a hurtful way all the while realizing that she was under tremendous stress and fully excusing her. In order to count as blaming for some behavior, it seems that the blamer must have certain beliefs about what the blamee did and whether she was culpable in doing it.

Buchak nicely draws our attention to the idea that blame seems inappropriate when we have only statistical evidence, and shows that something like (ST) helps vindicate this fact. Interestingly, her main purpose is to show that given that appropriate blame requires outright belief and that even very high credence is insufficient to justify blame, belief cannot simply be reduced to credence. (Not surprisingly, I'm very sympathetic to this conclusion and to her argument, and this is a point to which I will return briefly in section 5.) But, for present purposes, I note that the idea that our blaming practices do not seem to license blame solely on statistical evidence itself supports (ST) insofar as (ST) is part of the most natural explanation for this aspect of our practices of blame.

Because of the close connection between blame and punishment, it might be thought that we could take a simpler route to explaining the Proof Paradox than I offered in section 3. One might think that a guilty verdict simply expresses blame. But I do not think this accurately captures the relationship between blame and legal verdicts. Nevertheless, there is a connection. Arguably, *both* legal verdicts and blame, to be justified, depend on the rational belief that the agent in question is culpable. It is an open question whether the conditions for legal culpability include those for moral culpability, but in any case, belief in culpability of one kind or the other provides a common explanation for why neither is justified or rational based solely on statistical evidence. So, there is a strong connection between the two phenomena. At the same time, it is important to recognize that (ST) comes into play in explaining both legal practice and our

[33] This case overlaps with the bias cases presented in section 2. But we can imagine a similar case in which the features in question do not fit the typical stereotypical categories. For example, we could imagine finding out that people who like historical novels are ten times more likely to steal, and I think that this would not change our reaction to the case.

[34] See Brink and Nelkin (in preparation) and Clarke (2013) for some examples.

practices of blame, because these represent distinct, but central aspects of human life. And, as we will see, the differences between norms governing legal verdicts and interpersonal blame expand the resources for our defense of (ST).

5. Conclusion

Many have noticed that relying solely on statistical evidence raises special questions, and there is an on-going and serious debate about the propriety of P-inferences. We see parallels of this debate enacted in a variety of domains well beyond the Lottery Paradox and related paradoxes. In this paper, I have tried to bring these together to show that (ST), along with its rationale in terms of the nature of rationality, is part of a plausible explanation of our judgments and practices in a number of these domains.

Let us return very briefly to the deep concern about accepting (ST), namely, that it detaches rationality from accuracy when it should not, and to the extension of this concern that recommends moving away from the use of outright belief and toward the exclusive use of graded belief or credence. As we saw, there are at least two possible connections between rationality and accuracy. One is to see the value of accuracy as instantiated to the extent that true beliefs are acquired and fewer false ones are, even at a cost of arbitrariness. Douven (2008) is willing to accept arbitrariness for the sake of accuracy in his conception of rationality, or at least takes the burden of proof to be on those who assume otherwise (217). For example, he takes it that it is rational to believe of your ticket (t1) in a thousand ticket lottery that it will lose, and it is rational to believe the same of t2, and t3, and so on. But as we approach t1000, it would not be rational to believe that it will lose because at this point you would believe inconsistent things and you would thereby be licensed to add all kinds of false beliefs to your belief set. As long as you withhold belief about one ticket, you will have maximized your true beliefs at the risk of one false belief, which appears to be a good ratio (211). In contrast to this conception of rationality, I have offered one in which one may not achieve this sort of maximization of true beliefs even by being perfectly rational. (One will be entitled to a probabilistic belief about every ticket, but will not be entitled to outright belief about any ticket.) But one's beliefs will be based on evidence that is, in the key sense explained earlier, sensitive to the truth insofar as finding out that one's belief is false should lead one to question something about one's total package of reasons.

How to adjudicate between these two accounts, each of which sees rationality as a guide to truth? Some of the cases discussed earlier help illustrate at least one approach. On the view that accepts arbitrariness in rational belief in exchange for a premium on accuracy in the aggregate, we face a dilemma: either we reject any connection between rational belief and legal verdicts, or we end up with the unacceptable conclusion that as long as we stop convicting people after we reach a certain percentage of stadium event attendees (say, 70%) in Gatecrashers, we have not obviously done anything problematic. But clearly the arbitrariness of *action* that the arbitrary belief licenses is unacceptable. This suggests that something has gone wrong with our beliefs in the case, if indeed our verdicts are based on beliefs. Of course, one can take the first horn of the dilemma, denying that legal verdicts and rational belief are, or ought to be, closely connected, and there are resources in legal theory for making that case. But note that the

competing conception of rationality I offered has flexibility here to accept that rational belief is in fact required for coming to legal verdicts. And further, the particular resources available in legal theory for side-stepping rational belief are not available in the parallel practice of interpersonal blame. It would be odd to blame the first arbitrarily selected 70% of candidate gatecrashers, say, and then stop. And yet, we cannot appeal to special features of the function of the legal system to avoid this consequence of accepting arbitrariness. Perhaps the proponent of this sort of view can make the case for additional conditions on proper blame that go beyond those typically identified and that rule out justifiable blame in this case. But the case of blame poses a special challenge for this account.

Finally, consider the idea that we should relinquish the use of rational belief for any of these purposes, and instead simply rely on graded beliefs, or credences. This certainly provides an alternative picture. But as Buchak points out in the case of blame, and as we also see in the case of stereotypical beliefs, it appears that the distinction between outright belief and very high credence—or probabilistic belief—does a great deal of explanatory work. I believe that because of its explanatory power, the the package of commitments that together explain a variety of phenomena in the domains of biased belief, blame, and the proof paradox in the law provided here together constitute a significant reason for each individual commitment, including the retention of outright belief, the idea that actions are justified holistically by circumstances and sets of beliefs, among other things, and the solution to the Rationality Version of the Lottery Paradox by means of (ST).

References

Brink, David O. and Nelkin, Dana Kay (in preparation) "The Nature and Significance of Blame," *The Oxford Handbook of Moral Psychology* edited by John Doris and Manuel Vargas (Oxford: Oxford University Press).

Buchak, Lara (2014) "Belief, Credence, and Norms," *Philosophical Studies* 169 (285-311).

Christensen, David (2004) *Putting Logic In Its Place* (Oxford: Oxford University Press).

Cohen, L. Jonathan (1977) *The Probable and the Provable* (Oxford: Clarendon Press).

Douven, Igor (2008) "The Lottery Paradox and Our Epistemic Goal," *Pacific Philosophical Quarterly* 89 (204-225).

Enoch, David, Spectre, Levi, and Fisher, Talia (2012) "Statistical Evidence, Sensitivity, and the Legal Value of Knowledge," *Philosophy and Public Affairs* 40 (197-224).

Enoch, David and Fisher, Talia (2015) "Sense and Sensitivity: Epistemic and Instrumental Approaches to Statistical Evidence" *Stanford Law Review* (557-611).

Gardiner, Georgi (forthcoming) “Evidentialism and Moral Encroachment,” *Believing in Accordance with the Evidence: New Essays on Evidentialism*, edited by Kevin McCain (Springer).

Gendler, Tamar Szabó (2011) “On the Epistemic Cost of Implicit Bias,” *Philosophical Studies* 156 (33-63).

Harman, Gilbert (1968) "Knowledge, Inference, and Explanation." *American Philosophical Quarterly* 5 (164-73).

Ho, Hock Lai (2008) *A Philosophy of Evidence Law: Justice in the Search for Truth* (Oxford: Oxford University Press).

Kahneman, Daniel and Tversky, Amos (1996) “On the Reality of Cognitive Illusions,” *Psychological Review* 103 (582-591).

Kaplan, Mark (1996) *Decision Theory as Philosophy* (Cambridge: Cambridge University Press).

Kay, David (1978) “The Paradox of the Gatecrasher and Other Stories,” *Arizona State Law Journal* (101–43).

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

Moss, Sarah (2018a) “Moral Encroachment,” *Proceedings of the Aristotelian Society* CXVIII (177-205).

Moss, Sarah (2018b) *Probabilistic Knowledge* (Oxford: Oxford University Press).

Nelkin, Dana Kay (2000) “The Lottery Paradox, Knowledge, and Rationality,” *The Philosophical Review* 109 (373-409).

Nelkin, Dana Kay (in preparation) “Equal Opportunity: A Unifying Account of Moral, Epistemic, and Aesthetic Responsibility”.

Pereboom, Derk (forthcoming) “Free Will Skepticism and Crime Prevention” in *Free Will Skepticism in Law and Society* (Cambridge: Cambridge University Press).

Quillian, Lincoln and Pager, Devah (2010) “Estimating Risk: Stereotyple Amplification and the Perceived Risk of Criminal Victimization,” *Social Psychology Quarterly* 73: 79-104.

Redmayne, Mike (2008) “Exploring the Proof Paradoxes,” *Legal Theory* 14 (281-309).

Thomson, Judith Jarvis (1986) “Liability and Individualized Evidence,” *Law and Contemporary Problems* 49: 199-219.

Wheeler (2007) “A Review of the Lottery Paradox” *Probability and Inference: Essays in Honor of Henry E. Kyburg, Jr*. edited by William Harper and Gregory Wheeler (College Publications).