

2020

Reconsidering the rule of consideration: Probabilistic knowledge and legal proof

Timothy Smartt

The University of Notre Dame Australia, timothy.smartt@nd.edu.au

Follow this and additional works at: https://researchonline.nd.edu.au/ies_article



This article was originally published as:

Smartt, T. (2020). Reconsidering the rule of consideration: Probabilistic knowledge and legal proof. *Episteme, Early View Online First*.

Original article available here:

<https://doi.org/10.1017/epi.2020.28>

This article is posted on ResearchOnline@ND at
https://researchonline.nd.edu.au/ies_article/5. For more
information, please contact researchonline@nd.edu.au.



This is an Open Access article distributed in accordance with the Creative Commons Attribution Non-Commercial No Derivatives 4.0 International license (CC BY-NC-ND 4.0), which permits others to copy and redistribute the material in any medium or form, provided the original work is properly cited. See: <https://creativecommons.org/licenses/by-nc-nd/4.0/>

This is the author's version of the following article, as accepted for publication.

This article was originally published in *Episteme* available at:

<https://doi.org/10.1017/epi.2020.28>

Smartt, T. (2020) Reconsidering the rule of consideration: Probabilistic knowledge and legal proof. *Episteme, Early View Online First*. doi: 10.1017/epi.2020.28

Reconsidering the Rule of Consideration: Probabilistic Knowledge and Legal Proof *

Tim Smartt

timjsmartt@gmail.com

Forthcoming in *Episteme*

Abstract

In this paper, I provide an argument for rejecting Sarah Moss's recent account of legal proof. Moss's account is attractive in a number of ways. It provides a new version of a knowledge-based theory of legal proof that elegantly resolves a number of puzzles about mere statistical evidence in the law. Moreover, the account promises to have attractive implications for social and moral philosophy, in particular about the impermissibility of racial profiling and other harmful kinds of statistical generalisation. In this paper, I show that Moss's account of legal proof crucially depends on a moral norm called the *rule of consideration*. I argue that we have a number of reasons to be sceptical of this rule. Once we reject the rule, it is not clear that Moss's account of legal proof is either plausible or attractive.

Keywords: Probabilistic Knowledge, Mere Statistical Evidence, Legal Proof, Moral Encroachment, Sarah Moss.

*I'm grateful to Brian Hedden, Mark Colyvan, Daniel Greco, Caspar Hare, Kevin Dorst, an anonymous referee and associate editor of this journal, and an audience at MIT's SLLERG Seminar for very helpful discussion and comments that improved this paper. This research was supported by funding from the Society for Applied Philosophy.

This paper is about Sarah Moss's attempt to use new work in epistemology to solve an old problem in legal and social philosophy. The new work in epistemology is her theory of 'probabilistic knowledge' (Moss 2018a). The old problem in legal and social philosophy is the problem of mere statistical evidence. Her solution consists in defending a novel version of a knowledge-based account of legal proof, from which she also develops a unique account of what's wrong with beliefs about specific individuals based on social generalisations. I argue that her solution should ultimately be resisted.

This is newsworthy for three reasons. First, Moss's solution promises to deliver quite a few attractive benefits. A number of philosophers have acknowledged the benefits of a knowledge-based account of legal proof but have found that they come at too high a cost. Moss's new argument for this kind of account promises to deliver the benefits while avoiding the costs. Second, her solution claims to not just solve a problem in legal philosophy, but to have wider implications for moral and social philosophy; in particular, about the rational permissibility of statistical generalisations in social life, and the relationship between moral reasons and epistemic reasons. Third, the reason that her solution fails is instructive. It relies on a mechanism called 'the rule of consideration.' The arguments I offer against the rule of consideration show that any account of how we ought to accommodate probabilistic beliefs in legal and social life should do without mechanisms of this sort.

This paper is in four parts. In section 1, I motivate the problem mere statistical evidence presents for understanding legal proof. In section 2, I outline Moss's proposed solution and explain why it offers a very attractive way of thinking about knowledge, probability, and legal proof. In section 3, I introduce the rule of consideration, and outline the crucial role it plays in Moss's theory of legal proof. In section 4, I present two arguments for rejecting Moss's application of her epistemological view to the issue of legal proof. I also consider several ways the rule of consideration could be modified to avoid my challenges, and explore the costs and benefits of these possibilities. The central upshot of this paper is to show that in order for probabilistic knowledge to

generate attractive solutions to problems in legal and social philosophy it must rely on the rule of consideration, and we have good reasons to reject this rule as it's currently formulated.

1 The problem of mere statistical evidence

The problem of mere statistical evidence can be appreciated by considering another problem; namely, the problem of interpreting standards of legal proof. A simple and intuitive way of understanding the various standards of legal proof is in terms of probability.¹ According to this view, standards such as proof *by preponderance of the evidence*, proof *by clear and convincing evidence*, and proof *beyond a reasonable doubt*, establish probability thresholds which must be met for a verdict of guilt or liability. For instance, to find a defendant liable *by preponderance of the evidence*, it must be the case that the factfinder judges it to be more than 0.5 likely on the evidence that the defendant is liable. Similarly, to find a defendant liable *by clear and convincing evidence*, it must be the case that the factfinder judges it to be more than, say, 0.75 likely on the evidence that the defendant is liable.

Let's suppose that the legal probabilist intuition is more or less correct: standards of legal proof should be understood as picking our probability thresholds. The most serious problem with this view is that it seems that verdicts about individuals could be supported for each threshold merely on the basis of impersonal statistical evidence. This problem is famously illustrated in a set of puzzle cases.² Each of these toy cases describe a situation in which statistical evidence seems to establish a specific probability of guilt or liability, but we nevertheless feel reluctant to deliver such a verdict. The cases present a problem for legal philosophy broadly, since it is puzzling in a general sense why we might be reluctant to render a verdict of guilt or liability despite the presence of highly probabilifying evidence. However, the problem is amplified if one thinks that

¹For a recent articulation and defense of this view, see Colyvan and Hedden (2019).

²For helpful surveys, see Redmayne (2008) and Pardo (2019).

standards of proof *just are* probability thresholds. In this case, the puzzle cases are direct counterexamples. That is, they seem to show that verdicts of guilt or liability on any standard of proof must involve more than meeting a probability threshold, and thus standards of proof cannot be understood solely in terms of probability.

It will be worth looking briefly at some of the most prominent puzzle cases. PRISON YARD presents a case in the context of proof *beyond a reasonable doubt* (or something very close to it), and BLUE BUS in the context of a less demanding standard of proof, such as proof *by preponderance of the evidence*.

PRISON YARD: In an enclosed prison yard there are twenty-five identically dressed prisoners and a prison guard. The sole witness is too far away to distinguish individual features. He sees the guard, recognisable by his uniform, trip and fall, apparently knocking himself out. The prisoners huddle and argue. One breaks away from the others and goes to a shed in the corner of the yard to hide. The other twenty-four set upon the fallen guard and kill him. After the killing, the hidden prisoner emerges from the shed and mixes with the other prisoners. When the authorities later enter the yard, they find the dead guard and the twenty-five prisoners. Given these facts, twenty-four of the twenty-five are guilty of murder. There is a 0.96 probability, on the admitted facts, that any individual prisoner was involved in the murder. Thus, it seems that any prisoner can be found guilty of murder *beyond a reasonable doubt* (or something very close to it).³

BLUE BUS: A car is negligently run off the road by a blue bus. The driver of the car cannot identify the exact bus that caused the accident, but she can prove that the Blue Bus Company operates 80 percent of the blue buses in town, while another company operates only the remaining 20 percent. There is a 0.8 probability, on the admitted facts, that the bus in question was operated by the Blue Bus Company. Thus it seems that the Blue Bus

³PRISON YARD was originally presented in Nesson (1979).

Company can be found liable on a standard such as proof *by preponderance of the evidence*.⁴

Mere statistical evidence doesn't just present a problem for legal philosophy. It also generates confusion about the permissibility of inferences in broader social life based on statistical generalisations, especially inferences about individual people based on their gender, race, sexual orientation, religion, and so forth. The problem consists in the fact that our moral evaluation of such inferences tends to be negative, while our epistemic evaluation of the same inference can very often be positive. In this sense, statistical generalisations in social life present a dilemma between epistemic and moral concerns. This makes it unclear what we all-things-considered ought to do when in a position to draw an inference about an individual based on statistical evidence about members of their reference class. For instance, Tamar Szabó Gendler's COSMOS CLUB case illustrates the tension between epistemic and moral evaluations of these kinds of inferences.

COSMOS CLUB: On the night before he is to be presented with the Presidential Medal of Freedom, John Hope Franklin hosts a celebratory dinner party at the Cosmos Club. The Cosmos Club has very few African American members, and all the other African American men in the club that evening are uniformed attendants. While walking through the club, a woman sees him, calls him over, presents her coat check ticket and asks him to bring her coat.⁵

This case leaves us puzzled about what the woman ought to believe, since her epis-

⁴The original inspiration for BLUE BUS was *Smith v. Rapid Transit* 58 N.E.2d 754 (1945). For an early discussion of the implications of the case for statistical evidence and standards of legal proof, see Tribe (1971).

⁵Gendler (2011). Gendler's own view is that we ought to bite the bullet about the intractability of this dilemma: sometimes being rational will involve violating some moral norms. She writes, "In short, as long as there's a differential crime rate between racial groups, a perfectly rational decision maker will manifest different behaviors, explicit and implicit, towards members of different races. This is a profound cost: living in a society structured by race appears to make it impossible to be both rational and equitable." (Gendler 2011, 57). Another puzzle about the rational permissibility of inferences about individuals based on statistical evidence is *the reference class problem*. Individuals could be placed in very many different reference classes, which would justify a number of different, and in some cases jointly incoherent, inferences about them. I'll set this problem aside for the purposes of this paper. See, Colyvan, Regan, and Ferson (2001) and Hájek (2007).

temic reasons seem to pull her in one direction, and her moral reasons in another. It's worth noting here that the relevant question for our purposes is what the woman ought to believe, not how she ought to act. Were the puzzle about how the woman ought to act, the dilemma might easily be resolved; there is nothing especially puzzling about believing p but acting as if $\neg p$ in light of the expected utility of each option.

One way to solve the puzzle cases is to stipulate that mere statistical evidence is always insufficient for legal proof. A number of philosophers have been attracted to the idea that a knowledge-based account of legal proof provides a plausible solution along these lines.⁶ According to this view, legal proof requires the factfinder to know that the defendant is guilty or liable, and mere statistical evidence cannot lead to such knowledge. There are at least two strong motivations for the view that legal proof requires knowledge. First, the view provides a simple explanation of the inadequacy of mere statistical evidence for legal verdicts. Second, there seems to be a very natural connection between knowledge and legal proof. As Moss rightly emphasizes, when theorists attempt to define legal proof they often focus on properties that are widely taken to be hallmarks of knowledge. For instance, the following conditions have all been proposed as necessary for legal proof: sensitivity to the truth, incompatibility with luck, capable of serving as a reason for action, reliably safe from error, non-'Gettiered', and truth.⁷

However, despite these strong motivations, when one looks at specific standards of legal proof, it starts to seem implausible that proof by all standards requires knowledge. For instance, civil matters tend to require proof *by preponderance of the evidence*. It seems that proof *by preponderance of the evidence* can't require knowledge since it merely requires that the factfinder have a degree of confidence that is a long way from what's usually required for knowledge. Despite the many similarities between legal proof and knowledge, it looks as though we shouldn't help ourselves to the simple and ele-

⁶For an early exploration of the idea that legal proof requires knowledge, see Thomson (1986). For recent defenses, see Littlejohn (2017), Bloome-Tillmann (2017), and Smith (2018). For a helpful survey, see Gardiner (forthcoming).

⁷On sensitivity to truth, see Enoch, Spectre, and Fisher (2012). On incompatibility with luck, see Thomson (1986). On reasons for action, see Nesson (1985). On safety from error, see, Pritchard (2015). On legal proof and Gettier cases, see Pardo (2010). On truth, see Duff, Farmer, Marshall, and Tadros (2007). See also Moss's discussion of these views, Moss (2018a), 208–210.

gant explanation of these similarities; namely, legal proof in general requires knowledge. The incompatibility of knowledge with various standards of legal proof is significant enough that Mike Redmayne, in his survey article on standards of legal proof, sets aside knowledge-based accounts of legal proof solely because of this problem.⁸

We've seen that mere statistical evidence leaves us with a problem for understanding legal proof in terms of probability *and* a problem for understanding legal proof in terms of knowledge. For the legal probabilist, it seems that the mere statistical evidence featured in the puzzle cases provides a powerful counterexample to understanding standards of proof as probability thresholds. However, shifting to understanding standards of proof as requiring knowledge doesn't solve the problem. It seems that the salient standard in many legal proof contexts is much less demanding than knowledge and in these contexts the problem of mere statistical evidence remains. Moss's innovative account of legal proof offers a way to solve both problems at once.

2 Moss's knowledge-based account of legal proof

Moss's response to the problem of mere statistical evidence draws on her epistemological theory of 'probabilistic knowledge.' According to this theory, probabilistic beliefs, such as credences, can amount to knowledge in much the same way that outright beliefs can. For example, we can know that it's 0.6 likely that Amy is in Amsterdam, we can know that Beth is probably in her office, we can know that it's more likely than not that Clementine will win the election, and we can know that it might snow in London tomorrow. On Moss's view, probabilistic opinions like these can amount to, or fall short of, knowledge in all the familiar ways; in particular, by satisfying the most widely accepted necessary conditions for an opinion to count as knowledge, such as being safe,

⁸Redmayne writes: "There is an obvious problem with this view, however. It is plausible that whatever prevents a liability verdict in Prison Yard also prevents a liability verdict in [the civil case] Blue Bus. If Prison Yard is explained by a knowledge requirement for proof, then Blue Bus is too. But that would involve arguing that civil as well as criminal verdicts require knowledge, and that is not easy to accept. Civil verdicts require no more than proof on the balance of probabilities. This standard seems too low to satisfy the degree of justification required for knowledge." Redmayne (2008), 299.

non-‘Gettiered’, and true.⁹

Her account turns on a novel conception of mental and semantic content. Moss’s view is that many of our opinions have probabilistic content rather than propositional content. Probabilistic content is something of a term of art for Moss, and the formal details needn’t concern us too much for present purposes. However, it’s important to emphasise that in her view probabilistic contents don’t reduce to propositional contents. That is, when an agent knows that it might snow in London tomorrow, Moss would reject the notion that this attitude could only be a full belief in a proposition about probabilities, such as propositions about objective chances or evidential probabilities. Rather, Moss holds that opinions can be modelled using sets of probability spaces and these opinions can constitute knowledge.

Moss argues that one interesting implication of her theory of probabilistic knowledge is that it provides a new way of understanding legal proof in general as requiring knowledge. Moreover, the account solves the two problems outlined at the close of section 1: it provides a way for the legal probabilist to resist the notion that standards of proof can be satisfied by merely statistical evidence *and* it provides a knowledge-based account of legal proof that generalises across the various standards.

In particular, Moss’s notion of probabilistic knowledge is perfectly suited to solve Redmayne’s problem, which holds that knowledge-based accounts of legal proof cannot apply in weaker proof contexts. Probabilistic knowledge holds that probabilistic contents of various degrees of strength can constitute knowledge. On this view, then, there is no problem with treating weaker contents picked out by standards such as proof *by preponderance of the evidence* as candidates for knowledge. Moss uses this unique response to Redmayne’s problem to develop a new knowledge-based interpretation of the standards of legal proof. On her account, a defendant is proved liable *by preponderance of the evidence* only if the judge or jury have greater than 0.5 credence that the defendant is liable, and that credence constitutes knowledge. Likewise, a defendant is

⁹For a more detailed exposition of both Moss’s view and her arguments, see Smartt (2019).

proved liable *by clear and convincing evidence* only if the judge or jury knows an even stronger probabilistic content, and guilty *beyond a reasonable doubt* only if the judge or jury knows a still stronger probabilistic content.¹⁰ In general, we might say that a standard of proof is met when the judge or jury knows a specific content – namely, the set of probability spaces according to which the probability of guilt or liability meets or exceeds the relevant threshold for the standard. Moss’s proposal preserves the legal probabilist intuition about how to understand standards of legal proof. That is, she endorses the view that each standard corresponds to a probability threshold. Indeed, I think that one way to understand her view is as a specific version of legal probabilism, which adds a further necessary condition to any legal standard of proof: the judge or jury must know a probabilistic content which meets or exceeds the salient threshold. With this condition in place, one can avoid Redmayne’s problem that a probabilistic understanding of the various standards undermines understanding legal proof as requiring knowledge.

Furthermore, Moss argues that mere statistical evidence is usually insufficient for probabilistic knowledge.¹¹ Consider PRISON YARD. On her view, the mere statistical evidence in this case fails to provide the factfinder with probabilistic knowledge. Although it might justify a 0.96 credence in the claim that any individual prisoner committed the murder, this credence fails to be probabilistic knowledge. Thus you might have a very high justified credence that Smith (say) is probably guilty, but you do not *know* that Smith is probably guilty. This answer could then be filled out in a number of different ways, since there are very many familiar reasons why a justified belief might fail to constitute knowledge. However, the specific feature which Moss focuses on is the condition that one does not know p if one is not in a position to rule out a contextually relevant alternative to p . In the case of PRISON YARD, a factfinder cannot *know* that Smith is probably guilty, since the mere statistical evidence does not allow her to rule out a relevant alternative; namely, that Smith is the one exception within the reference

¹⁰Picking out a precise and invariant threshold value for the *beyond a reasonable doubt* standard can seem arbitrary, and Moss notes that her view is compatible with holding that this value is context-sensitive. Moss (2018a), 212.

¹¹As we’ll see in section 4, Moss makes an important concession in the case of BLUE BUS.

class. Moss argues that not only does this solution show why the evidence in PRISON YARD fails to provide probabilistic knowledge that Smith is probably guilty, but it can be extended to an explanation of the insufficiency of mere statistical evidence in other proof contexts. Although the strength of the probabilistic content will vary as the standard of proof varies, the basic conditions on probabilistic knowledge remain fixed. So even in proof contexts involving much lower probabilistic thresholds, the mere statistical evidence might provide justification for a certain credence but not knowledge of this content.

One attractive benefit of this solution to the problem of mere statistical evidence is that it generalises beyond legal philosophy. For example, Moss treats statistical generalisation in social life in a similar way. In cases of social generalisation where our evidence consists only of mere statistical evidence, we are unable to rule out the possibility that an individual person is unlike an arbitrary member of their reference class. Thus we will not be in a position to *know* that the individual probably has the target property.

3 The rule of consideration

Moss's application of probabilistic knowledge to legal proof generates an interesting puzzle. The puzzle consists in the fact that, on Moss's account, probabilistic knowledge of lottery propositions and probabilistic knowledge in legal and social contexts are surprisingly disanalogous.¹² In PRISON YARD, Moss holds that we cannot know that Smith is probably a murderer, since we are not in a position to rule out the possibility that Smith is an unrepresentative member of his reference class. By applying the same reasoning to a lottery case, we meet the result that we cannot know that our single ticket in a large lottery is probably a loser, since we are ordinarily not in a position to rule out the possibility that our ticket is an unrepresentative member of its reference class,

¹²This puzzle is not acknowledged in Moss (2018a), Moss (2018b), or Moss (forthcoming). Moss does address some analogies between lottery knowledge and knowledge in legal and social contexts, but doesn't address the disanalogy between probabilistic lottery knowledge and probabilistic knowledge in legal and social contexts. See, for example, Moss (2018a), 216–220, Moss (2018b), 187–188, and Moss (forthcoming) §1 and §3.4.

namely, the winning ticket. But this seems absurd! Knowing that one's lottery ticket is probably a loser is a paradigmatic instance of probabilistic knowledge. So it seems that the approach to probabilistic knowledge in PRISON YARD overgeneralises and rules out probabilistic knowledge of lottery propositions.

The force of the puzzle can also be felt in the opposite direction. That is, if Moss *were* to treat probabilistic knowledge of lottery propositions and probabilistic knowledge in legal and social contexts analogously, this would lead to some very problematic results for her account of legal proof. Suppose we can know that our lottery ticket is probably a loser. By analogy, it would seem that we could sometimes come to know contents about individuals based on mere statistical evidence or statistical generalisations. That is, if we treat probabilistic knowledge of lottery propositions and probabilistic knowledge in legal and social contexts analogously, we meet the unhappy results that we can *know* that Smith is probably one of the murderers, we can *know* that Franklin is probably a waiter, and so on. This would strike many philosophers as regress, not progress, in social and legal philosophy. In any case, it's exactly the opposite lesson Moss draws.

Perhaps to avoid this puzzle, Moss introduces an important piece of machinery – called the *rule of consideration* – which will be the focus of my objections in section 4. The rule states,

Rule of consideration: In many situations where you are forming beliefs about a person, you morally should keep in mind the possibility that they may be an exception to statistical generalizations. (Moss 2018a, 221)

Since the rule only applies to people and not objects, we are left with a fundamental distinction between probabilistic knowledge of people and probabilistic knowledge of objects (Moss 2018a, 218, 219, 223). This norm allows you to know that your lottery ticket – an object – is probably one of the losing tickets, based only on the statistics about members of the salient reference class. But the norm prevents you from knowing that Smith – a person – is probably one of the murderers, based only on the statistics about members of the salient reference class.

Moss also relies on the rule of consideration in her application of probabilistic knowledge to social philosophy. She argues that in cases involving potential probabilistic knowledge of people we are morally obligated to consider the alternative that an individual person is an exception to the statistics. That is, in cases of profiling and social generalisation we are morally obligated to consider alternatives which will block probabilistic knowledge. This leads Moss to develop a hybrid norm against profiling, which spells out our epistemic obligations in contexts where the rule of consideration applies. The hybrid norm holds that it morally ought to be the case that it epistemically should not be the case that you form beliefs about people based on mere statistical evidence. (Moss 2018a, 222ff).

If we consider Gendler's COSMOS CASE, this application has the attractive feature of dissolving the tension between moral and epistemic evaluations of the woman's inference. It seemed that there's a moral case for thinking the woman judged as she oughtn't to have and an epistemological case for thinking she judged as she ought to have. However, Moss's hybrid norm provides us with the result that had the woman been acting as she morally ought to act, then it would have been the case that she epistemically ought not to form this judgment. Thus, on Moss's account, there is a sense in which refusing to draw inferences about individuals based on statistical generalisation is not making an epistemic compromise. Rather, it is judging as one epistemically ought to in cases in which the moral rule of consideration applies.

4 Objections to the rule of consideration

In this section I raise two separate objections to Moss's distinction between probabilistic knowledge of objects and persons. The objections motivate my conclusion that we should reject the rule of consideration – at least as it's currently formulated. This is significant since Moss relies on the rule to both avoid unattractive results about probabilistic knowledge and to build her positive case for the useful applications of probabilistic knowledge

to legal and social philosophy. In light of these objections, I think we should resist Moss's claim about the legal and social applications of probabilistic knowledge.¹³

4.1 Applying the rule

My first objection is that there are cases in which it is unclear how we should apply the rule; not because it's unclear whether a specific item is a person or an object, but because it's unclear what you're fundamentally forming your credence about. I think BLUE BUS provides us with such an example. In this case, it's usually stipulated that we're forming credences about buses; that is, based on the statistical evidence of the case we can have a 0.8 credence that the bus which caused the accident was operated by the Blue Bus Company. But there are at least two other plausible interpretations of the object of our credence in this case. First, our credence might be about the Blue Bus Company. That is, based on the statistical evidence of the case, we might have a 0.8 credence that the Blue Bus Company is the company responsible for the accident. Second, our credence might be about a bus driver (every bus in BLUE BUS has a bus driver, after all). That is, we might have a 0.8 credence that the driver who caused the accident works for the Blue Bus Company. So in BLUE BUS, our attitude could be interpreted as about an object, a group agent, a person, or perhaps a combination of all three.

In Moss's own discussion of BLUE BUS it actually isn't clear whether the rule applies to this case. Since the case is usually interpreted as involving credences about buses, not people, it seems that the rule of consideration doesn't apply in this case, and so Moss entertains the idea that perhaps mere statistical evidence *is* enough to provide

¹³Moss (2018b) provides further discussion of her theory of legal proof, the harms of profiling, the nature of moral stakes, and the rule of consideration. Importantly for my purposes, Moss (2018b) makes a slight revision to her position on the rule of consideration (see 191–192). The revision is that Moss states that she does not believe that the epistemic impact of the rule of consideration is an instance of moral encroachment (191). However, despite this clarification Moss (2018b) stands by the features of the rule that motivate my criticism. In particular, she continues to assign the rule a substantial role in her account of legal proof and social generalisations, and she continues to draw a fundamental distinction between probabilistic knowledge of objects and probabilistic knowledge of people. Clarifying whether or not the rule of consideration is an instance of moral encroachment is a welcome refinement to Moss's overall position, but this does not impact the arguments presented in this paper for rejecting Moss's account of legal proof.

probabilistic knowledge in BLUE BUS. However, she resists taking a firm stand about the case. Moss does note that BLUE BUS tends to generate more diverse intuitions than PRISON YARD. In particular, she writes that whilst verdicts of guilt or liability based on mere statistical evidence involving people seem straightforwardly intolerable, “courts are sometimes willing to act as if statistical inferences about *objects* can ground knowledge.” (Moss 2018a, 219). I find it interesting that Moss resists making a judgment about BLUE BUS. One explanation for this hesitancy might be that it is not clear how we ought to apply the rule in this case, and so it is not clear whether the statistical evidence in the case can give rise to probabilistic knowledge. In PRISON YARD, our credences are unambiguously about people so the rule applies. But in BLUE BUS, our credences are open for interpretation. In the remark quoted above, it seems that Moss treats them as being about objects. If this is true, then the rule needn’t apply, and BLUE BUS is analogous to a lottery case, and there is no principled reason why we cannot have probabilistic knowledge in this case. That Moss declines to draw this conclusion could be explained in many ways. Nevertheless, for our purposes I take it that one explanation of her reluctance to make a judgment about BLUE BUS is that it is unclear how to apply the rule to this case.

So far, I’ve emphasised that it’s unclear how to apply the rule when one’s credence isn’t paradigmatically about an object or a person. There are a number of ways that credences might have this feature, and attempting to apply the rule to our credences in BLUE BUS helpfully illuminates at least two of these. First, sometimes attitudes are about neither a person nor an object, for instance, they can be about a group agent. Importantly for Moss’s purposes, attitudes about group agents don’t just feature in fanciful thought experiments – they often occur in actual contexts of social and legal importance. For instance, in civil cases like BLUE BUS we’re interested in whether the factfinder knows that the *defendant* is probably liable, and defendants can often be corporate entities. I take it that Moss would want some version of the rule to apply to cases involving group agents, otherwise her account would permit findings of guilt or

liability about group agents on the basis of mere statistical evidence. Nevertheless, as the rule currently stands, it's not clear whether the rule applies in cases like this. Second, sometimes our credences don't refer *de re* to a particular entity, but they do refer to an entity under the guise of a definite description. For instance, based on the statistical evidence available in BLUE BUS, we might have a 0.8 credence that the driver who caused the accident works for the Blue Bus Company. But in this case, it's not clear whether the rule applies. Perhaps it does, since your attitude is about *some* person (although you do not know who this is). But perhaps it doesn't, since your attitude doesn't make a statistical inference about any particular person. Plenty of attitudes in weighty social and legal contexts can have this feature. Here's one motivation for thinking the rule should apply in cases like this. Suppose I know very little about politics in New Zealand, but – through a few occasional glances at international news – I do know the names of some current politicians and the statistics about the gender distribution of all its Prime Ministers. It's not clear that the rule applies if I form a high credence that the Prime Minister of New Zealand is probably a man, based just on the gender statistics. But it is clear that the rule applies (and that I violate it) if I form a high credence that Jacinda Ardern is probably not the Prime Minister of New Zealand, based just on the same statistics. In light of this case, it seems to me that we have some reason to think that the rule should apply to both attitudes, since the moral and epistemological differences between them are fairly minimal. At the least, it is unsatisfying that it is not clear whether the rule applies in cases like this.¹⁴

We've seen that trying to apply the rule to BLUE BUS raises a number of complex issues in respect to how our credences might fail to be paradigmatically about an object or a person. Let's set aside this particular case and consider a cleaner example which captures the central worry I wish to raise; namely, that in many cases it's unclear whether the rule applies. Imagine the following case involving three rooms.¹⁵ The first room contains 100 hats. The second room contains 100 people. The third room contains

¹⁴Thanks to a referee for helpful comments on this point.

¹⁵Thanks to a referee for suggesting this particular example.

100 people, each of whom is wearing a hat. I take it that Moss would hold that the rule of consideration applies when considering whether opinions about the people-only room count as knowledge and that it doesn't apply when considering opinions about the hats-only room. Does it apply when considering whether opinions about the people-wearing-hats room count as knowledge? If it doesn't, this seems unmotivated, since it's not clear whether we should interpret credences about the entities in this room as about objects or people. If it does, this leads to counterintuitive results. Suppose that in both rooms containing hats, 96 hats are red and 4 hats are blue. If the rule of consideration applies when forming opinions about the people-wearing-hats room, we get the result that we can know, when conditions are right, that a randomly selected hat from the hats-only room is probably red, but we fail to know, under the same conditions, that a randomly selected person wearing a hat from the people-wearing-hats room is probably wearing a red hat. At best, it seems unclear whether we should apply the rule in the people-wearing-hats case. At worst, applying the rule here seems to provide strange results when compared with nearby cases involving simply people or simply objects.¹⁶

Let's take stock. I've outlined a general challenge that prevents us from clearly applying the rule of consideration: sometimes our credences aren't paradigmatically about a person or an object. Sometimes they could be interpreted either way, as in the case of an opinion about people wearing hats. Sometimes they're about something which is neither a person nor an object, as in the case of an opinion about a group agent. And sometimes they can be about a person or an object in such a way that they fail to pick out an individual, as in the case of an opinion about an entity under the guise of a definite description.

How might Moss reply to this challenge? One option would be to broaden the rule such that it applies when an opinion is at least *partly* about a person.¹⁷ This would

¹⁶Thanks to Daniel Greco and Caspar Hare for helpful comments on this point.

¹⁷As a referee helpfully pointed out, recent work on the topic of 'aboutness' in the philosophy of language has provided new motivations for taking seriously the idea that sentences can be partly about more than one subject matter at once, that sentences can be partly true, and that sentences can be partly or entirely about subject matters that they do not explicitly mention. See Yablo (2014).

provide the result that the rule clearly applies in cases involving people-object hybrids, and in cases involving group agents, and would provide some motivation to think it also applies in cases involving people under a definite description. So it looks like modifying the rule along these lines would avoid the challenge I've outlined. But it also brings its own costs. In particular, this modification requires care to ensure that it doesn't overgeneralise. For instance, some of the canonical lottery cases are partly about people: people having surprise heart attacks, or people stealing cars, or people getting holes-in-one.¹⁸ Were the rule to apply when an attitude is partly about a person, this would make it difficult to have probabilistic knowledge in these kinds of cases. This would be an unattractive result for at least two reasons. First, as noted in section 3, it seems that probabilistic lottery knowledge – knowing that one's lottery ticket is probably a losing ticket – is a paradigmatic instance of probabilistic knowledge. Second, it's plausible that whatever we say about lottery knowledge should generalise to other cases with the same structure. A modified rule would introduce a seemingly arbitrary distinction into the kinds of lottery cases that are compatible with probabilistic knowledge. This division of lottery cases would provide the results that, for example, we often can know that our lottery ticket is probably a loser, but we often cannot know that all sixty golfers in a tournament will probably not score holes-in-one on a single hole.¹⁹

¹⁸See Hawthorne (2004), chapter 1. As McGrath (2004) puts it, "In effect, we enter a lottery when we leave our parked cars (winners have their cars stolen), and even by being just by being alive (winners die of a heart attack next year, or tomorrow)."

¹⁹The hole-in-one example is HEARTBREAKER: "Sixty golfers are entered in the Wealth and Privilege International Tournament. The course has a short but difficult hole, known as 'Heartbreaker.' Before the round begins, you think to yourself that, surely not all sixty players will get a hole-in-one on the 'Heartbreaker.'" (See Hawthorne (2004), 12). I take it that the rule of consideration would apply to forming opinions about golfers in HEARTBREAKER based on statistical generalisations about how golfers perform on the 'Heartbreaker', which would prevent you from knowing of any given golfer that they probably won't get a hole-in-one. The general point is that when a case includes (or entails) a lottery proposition that is about (or partly about) people, such as golfers, the rule of consideration will often block probabilistic knowledge, whereas when a case includes (or entails) a lottery propositions that is only about objects, such as lottery tickets, the rule of consideration will not apply and probabilistic knowledge will often be possible.

4.2 Justifying the rule

We’ve seen that one possible way for Moss to avoid my first challenge is to stipulate that the rule applies whenever a credence is at least partly about a person. We’ve also seen that this threatens to overgeneralise, making it difficult to have probabilistic knowledge in a range of cases involving people. One way to avoid this overgeneralisation would be to *narrow* the rule along a different dimension, such that it only applies when you risk harming a person by forming a false belief about them based on statistical generalisations.²⁰ This would provide a principled way of ensuring that the rule doesn’t apply when an opinion is about (or partly about) a person but does not expose them to any moral harm.²¹ This raises a second challenge that I’ll focus on in this section. The challenge is that it’s not clear what justifies a moral distinction between probabilistic beliefs about objects and people. Moss says little about what justifies the rule, but everything she does say has to do either with stakes or with concerns about respecting someone as an individual. I’ll argue that neither justification grounds the moral distinction made by the rule.

Let’s take stakes first. Moss writes that whether we should form a belief based on

²⁰As a referee helpfully pointed out, if you think that statistical generalisations can be morally objectionable on grounds apart from risks of harm, then you may find this position unsatisfying. For example, consider Rima Basu’s discussion of the RACIST HERMIT case. This case involves a hermit who forms a belief about someone with whom they will never interact on the basis of a statistical generalisation. The case stipulates that the belief has pejorative racial content, is never expressed, and is true. Basu holds that such a belief still constitutes a moral wrongdoing. Those who share Basu’s intuition about RACIST HERMIT might hold that the narrowed version of the rule of consideration discussed in this section fails to capture the moral problem with statistical generalisations. See Basu (2019a), 919 and Basu (2019b), 2504.

²¹Some philosophers have assumed that the rule should be interpreted this way. For instance, Daniel Greco, in the context of discussing the rule of consideration, writes, “I assume I can know that someone probably speaks some French, given that they grew up in Montreal. While they *might* be an exception to the generalization that most people who grew up in Montreal speak some French, I assume there’s no moral problem with failing to keep exceptions to that generalization in mind – i.e. failing to treat it as a relevant alternative that must be ruled out if I am to have probabilistic knowledge – in attributing to myself or others knowledge that they speak some French.” Greco notes that, “we’d like some story about just what sorts of generalizations about people are ones whose exceptions must be kept in mind, from a moral point of view.” Greco (2020), 114. Moss (2018b) gives us reason to think that she would be sympathetic to this modification to the rule. She writes that “Whether statistical evidence is sufficient for knowledge depends partly on what is at stake.” She claims that some kinds of profiling and opinions based on statistical generalisations might be beneficial, and these opinions needn’t be subject to moral encroachment. To support this, she provides an example involving the beneficial use of racial profiling by medical experts when forming an opinion about the likelihood that an individual has a disease on the basis of statistics about their racial group. See Moss (2018b), 200. See also Moss (2018a), 223–224.

statistical inference “may depend partly on what is at stake if your belief turns out to be false.” (Moss 2018a, 223). I’m inclined to agree. However worries about the stakes involved in false probabilistic beliefs can be extended to objects too. For instance, take a probabilistic belief about my single ticket in a billion-dollar lottery. It seems that I can know that my ticket will probably lose the lottery. But this is a high-stakes belief – I stand to lose a billion dollars if I’m wrong!

Perhaps Moss might reply that it is not stakes in general that we are interested in, but moral stakes in particular, understood in terms of false beliefs which unacceptably run the risk of causing harm to real people.²² However, it still seems that probabilistic opinions about objects can have high moral stakes too. That is, I think a false probabilistic belief about an object could harm a real person who is not themselves the object of the belief. For instance, consider a probabilistic belief about the health of a fish I just caught in Sydney Harbour. Suppose it’s the case that 96% of the fish in Sydney Harbour are healthy and suitable for human consumption. On Moss’s view, if conditions are right, I can know that this fish is probably healthy. However, there are moral stakes if I’m wrong! If I’m wrong and I serve the fish to guests at my dinner party, I could cause them all harm in the form of awful food poisoning. A number of probabilistic beliefs about objects might serve as premises in practical reasoning in similar ways which risk harming people if they are false.

Furthermore, we can think of an example where it’s initially not clear whether the rule applies, because our attitude isn’t paradigmatically about a person or an object, but were the moral stakes of the case to be raised, then it starts to seem like the rule should apply. For instance, consider my example of the three rooms containing people, hats, and people wearing hats. The upshot of this case was that it’s unclear whether the rule applies to opinions about the people-object hybrids in the people-wearing-hats room. Let’s suppose that in the previous case the people wearing red hats did so as a mere fashion choice. Let’s now suppose that the red hats are official merchandise

²²Moss (2018b) makes it clear that she is primarily interested in the moral stakes associated with probabilistic opinions, and less concerned with practical stakes in general.

endorsing a controversial politician. In this modified case, the people wearing red hats do so as an act of political expression. Let's hold all the other details fixed. There's still 96 red hats and 4 blue hats in each room that contains hats; it's just that now the red ones also bear a political slogan. I take it that the rule still doesn't apply when forming opinions about the hats-only room; that is, we can know, when conditions are right, that a randomly selected hat from this room is probably a red hat. But the people-wearing-hats room now carries some moral stakes – by forming an opinion about whether a randomly selected person from this room is probably wearing a red hat based on the statistics about red hats in the room, we run the moral risk of falsely attributing controversial political sympathies to her. So it seems like the rule should apply. But notice that it's the moral stakes that activate the rule, as it were. It's no clearer that our opinions in the high-stakes people-wearing-hats room are paradigmatically about people compared with our opinions in the low-stakes equivalent. What has changed between the two cases, is that our opinions in the high-stakes case now run moral risks which they didn't in the low-stakes case.²³

The lesson is that moral stakes cannot justify the rule's distinction between probabilistic knowledge of objects and persons since they don't only attach to opinions about persons. Moral stakes can sometimes attach to opinions about objects or opinions that are neither paradigmatically about objects or persons.

I take it that one possible line of response to this problem will not be attractive to Moss. That is, she could broaden the rule of consideration such that it requires that one consider the possibility that a member of a reference class is unrepresentative – be it an object, a person, or whatever – when one's probabilistic beliefs about that member generate moral risks for real people. This way of modifying the rule would provide the results that I don't know that the fish is probably healthy and that I don't know that a randomly selected person from the high-stakes people-wearing-hats room is probably wearing a red hat. However, I take it that this reformulation would not be attractive

²³Thanks to an associate editor for helpful comments on this point.

to Moss as it would create problems for her account of legal proof. Plenty of legal evidence is probabilistic in nature.²⁴ For instance, many different kinds of evidence drawn from methods such as DNA testing, psychological evaluations, eyewitness testimony, or forensic accounting can be associated with statistical information about the likelihood of errors involved in each method. Since this evidence can feature in legal judgments that run the moral risk of harming real people with inappropriate verdicts and punishments, a much wider rule of consideration would require the factfinder to consider the relevant alternative that each piece of evidence is unlike typical members of its reference class. Were this the case, probabilistic knowledge would be much harder to have in legal contexts, as the factfinder would often not be in a position to rule out these relevant alternatives. Even in contexts involving a low standard of proof, such as proof by *preponderance of the evidence*, implementing this broadened rule would have the consequence that one often fails to *know* the relevant content necessary for legal proof. For example, imagine a factfinder considering whether it is at least 0.501 likely that Amy is liable for an offense. The main evidence in the case is CCTV footage that seems to show Amy committing the offense. Amy stands to be harmed by a false verdict, so the factfinder ought to consider the possibility that the CCTV evidence is unlike most other CCTV evidence – which is highly reliable – and has, say, been tampered with by Amy’s arch-nemesis. Where the factfinder is not in a position to rule out this relevant alternative, the factfinder would fail to know that it’s at least 0.501 likely that Amy is liable, and thus Amy could not be found liable on even this low standard of proof. So I take it that broadening the rule to apply to objects which pose moral risks to people will not be attractive to Moss.

So far we’ve seen that doubts can be raised about justifying the rule in terms of stakes. I’ve provided reason to think that moral stakes can attach to probabilistic beliefs about objects and people-object hybrids. For the reasons mentioned above, I suspect that

²⁴Some theorists think that *all* evidence is probabilistic in nature. For instance, Judge Posner writes, “All evidence is probabilistic, and therefore uncertain; eyewitness testimony and other forms of ‘direct’ evidence have no categorical epistemological claim to precedence over circumstantial or even explicitly statistical evidence.” See *Milan v. State Farm Mutal Automotive Insurance Co.* 927 F.2d 166, 170 (1992).

Moss will not want the rule to extend to high-stakes probabilistic beliefs about objects. The upshot is that it seems that a moral distinction between probabilistic knowledge of persons and objects cannot be justified in terms of the moral stakes that accompany each kind of attitude.

Here's another possible justification for the rule that Moss seems to be sympathetic towards: we are morally required to respect individual persons in a way that we are not morally required to respect individual objects. That is, I suspect Moss could reply to my challenges so far by asserting that we really do owe individual people respect in a way that is different to the respect we owe to individual objects, even if symmetrical moral risks attach to a token opinion about an object and a token opinion about a person. But this justification makes the rule quite a controversial piece of normative ethical theory. I take it that Moss will want her account of legal proof to be compatible with many different views about, for instance, who is owed respect and why. Were the rule of consideration to involve substantial moral commitments, its appeal – and, in turn, the appeal of her account of legal proof – would be narrowed.

There's no need to embrace too sceptical a conclusion at this point. Perhaps it's possible to justify the rule in moral considerations about respect for persons in a way that avoids taking on divisive commitments in normative ethics.²⁵ However, were Moss to attempt to justify the rule in this way, I think she would face at least two challenges.

The first challenge is that Moss will not be able to straightforwardly help herself to one of the most prominent accounts of respect; namely, Stephen Darwall's account of 'recognition respect' and 'appraisal respect' (Darwall 1977). Like Moss, Darwall is interested in the kinds of respect morally owed to persons and objects. But he marks this distinction quite differently. According to Darwall, we can owe 'recognition respect' to both persons and non-persons, including objects, institutions, and non-human parts

²⁵For instance, I think Moss could avoid the possible challenge that one could only justify the rule in moral considerations about respect for persons by accepting a deontological moral framework. There are good reasons to think that a respect for persons norm is available to a range of normative ethical theories. For an argument that a fundamental moral requirement to respect persons is compatible with a consequentialist moral framework, see Pettit (1989).

of the natural world. ‘Recognition respect’ requires that we give appropriate weight in our practical reasoning to some fact about the entity that generates this kind of respect, and to then act appropriately. Sometimes ‘recognition respect’ generates distinctly moral requirements. That is, sometimes respecting something in this sense will involve taking into account some feature of the entity which places restrictions on what actions are morally permissible. For example, on this sense of respect, the fact that a species of flower is endangered can generate respect for the flower that constrains how we ought to act towards it. ‘Appraisal respect’ has its exclusive object as persons, and consists in the positive evaluation of a person’s actions or character. Darwall emphasizes that this sense of respect doesn’t place moral constraints on our behaviour. Respecting a person in this sense might generate prudential requirements or it might entail that we ought to praise or admire the person, but ‘it doesn’t essentially involve any conception of how one’s behavior towards that person is appropriately restricted’ (Darwall 1977, 41). For instance, if we respect a brilliant novelist or inspiring leader in this sense, it’s not the case that we’re under any special moral requirements about how we ought to treat them. It seems to me that Darwall’s account of respect will not help justify the rule of consideration. The kind of respect that most closely resembles the rule’s moral concern for persons is ‘recognition respect,’ but this is the kind that Darwall believes can be owed to both persons and non-persons. The kind of respect that Darwall believes is exclusively owed to persons doesn’t generate moral requirements like Moss needs and, in any case, it doesn’t capture the sense in which the rule is concerned about respecting persons.

Let’s set aside Darwall’s characterisation of different kinds of respect. Perhaps the rule could be justified by the notion that we should treat people as individuals.²⁶ This strategy raises the second challenge; namely, that the requirement to treat a person as an

²⁶At some points, Moss suggests that the rule is justified in this way. For example, she says, “The rule of consideration spells out one modest interpretation of the thought that we should *treat people as individuals*...As I see it, there are actually several moral norms corresponding to the rough idea that people should be treated as individuals, and the rule of consideration is among the least demanding of these norms.” (Moss 2018a, 223, emphasis hers).

individual needn't support the rule of consideration. The rule requires that one keep in the mind the possibility that a person is an exception to the generalisation. But I think it's plausible that one can treat a person as an individual without keeping this possibility in mind. Let's suppose that respecting a person as an individual involves being open to the possibility that the individual is an exception to the statistical generalisation. There are at least two plausible senses in which we might 'remain open' to this possibility. The first is *synchronic*: at a time, you do not rule out the possibility that the individual is an exception to the generalisation. This sense motivates the rule of consideration. The second is *diachronic*: at a time, you *do* rule out the possibility that the person is an exception, but you diachronically remain open to the possibility that the individual is an exception by being disposed to update your attitude in light of new evidence. The diachronic sense of 'remaining open' to the possibility that an individual is an exception is compatible with violating the rule of consideration at a particular time-slice. I take it that the diachronic sense of 'remaining open' is one plausible sense of the idea. You would *not* remain open were you disposed to dismiss or discount evidence about the individual, neither of which are true in the diachronic case. So it looks like the bare notion of treating a person as an individual is not enough to justify the rule.²⁷ One option would be to attempt to justify the rule in a more specific notion of what is involved in treating someone as an individual. For example, Eidelson (2013) and Wasserman (1991) argue that the law ought to treat people as individuals *by* respecting their autonomous agency. However, the more specific we get about what is involved in treating a person as an individual, the harder it becomes to avoid controversial moral commitments.

I've argued that that we can raise challenges for justifying the rule in either of the two

²⁷An example of some philosophers who believe that treating a person as an individual is compatible with violating the rule of consideration is Di Bello and O'Neil (2020). They write, "But treating someone as an individual does not, on the most plausible interpretation, require eschewing reliance on (accurate) generalizations. Rather, it only requires being receptive to the possibility that the individual is an exception to the generalization—that is, being receptive to any finer-grained information that would indicate that the generalization is not to be relied on in drawing a conclusion about the individual. . . . But a defendant in a trial is given an opportunity to introduce evidence that distinguishes them from their group and is entitled to have this information taken into consideration by the fact finders. In a well-functioning courtroom, then, fact finders could use profile evidence without failing to treat the defendant as an individual." (Di Bello and O'Neil 2020, 154).

ways that Moss is sympathetic towards; namely, in terms of moral stakes or in terms of moral considerations about respecting individual persons. In light of this, it might be the case that Moss needs to bite a bullet and justify the rule by accepting some substantial moral commitments. At a minimum, it is not clear whether the rule can be justified in a way that leaves it as morally uncontroversial as Moss takes it to be.

5 Conclusion

The advertised benefits of Moss's innovative knowledge-based account of legal proof are very attractive. I've argued that a crucial part of delivering these benefits is constraining probabilistic knowledge in legal and social life with the rule of consideration, and I have raised two separate challenges to the rule. For the legal and social benefits of probabilistic knowledge to be secured, Moss owes us a deeper account of how to apply and how to justify the rule of consideration.

References

- [1] Basu, R. (2019a). What we epistemically owe to each other. *Philosophical Studies*, 176, 915–931.
- [2] Basu, R. (2019b). The wrongs of racist beliefs. *Philosophical Studies*, 176, 2497–2515.
- [3] Bloome-Tillmann, M. (2017). 'More likely than not': knowledge first and the role of bare statistical evidence in courts of law. In Carter, J. A., Gordon, E. C., & Jarvis, B. (Eds.) *Knowledge First: Approaches in Epistemology and Mind* (pp. 278–292). Oxford: Oxford University Press.
- [4] Colyvan, M., Regan, H. M., & Ferson, S. (2001). Is it a crime to belong to a reference class? *Journal of Political Philosophy*, 9(2), 168–181.
- [5] Colyvan, M. & Hedden, B. (2019). Legal probabilism: a qualified defence. *Journal of Political Philosophy*, 27(4), 448–468.

- [6] Darwall, S. (1977). Two kinds of respect. *Ethics*, 88(1), 36–49.
- [7] Di Bello, M., & O’Neil, C. (2020). Profile evidence, fairness, and the risks of mistaken convictions. *Ethics*, 130(2), 147–178.
- [8] Duff, A., Farmer, L., Marshall, S., & Tadros, V. (2007). *The Trial on Trial (vol. 3): Towards a Normative Theory of the Criminal Trial*. Oxford: Hart Publishing.
- [9] Eidelson, B. (2013). Treating people as individuals. In Hellman, D., & Moreau, S. (Eds.) *Philosophical Foundations of Discrimination Law* (pp. 203–227). Oxford: Oxford University Press.
- [10] Enoch, D., Spectre, L. & Fisher, T. (2012). Statistical evidence, sensitivity, and the legal value of knowledge. *Philosophy and Public Affairs*, 40(3), 197–224.
- [11] Gardiner, G. (forthcoming). Legal evidence and knowledge. In Lasonen-Aarnio, M., & Littlejohn, C. *The Routledge Handbook of the Philosophy of Evidence*. New York: Routledge.
- [12] Gendler, T. (2011). On the epistemic costs of implicit bias. *Philosophical Studies*, 156(1), 33–63.
- [13] Greco, D. (2020). Acting on probabilistic knowledge. *Res Philosophica*, 97(1), 109–117.
- [14] Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
- [15] Hajek, A. (2007). The reference class problem is your problem too. *Synthese*, 156(3), 563–585.
- [16] Littlejohn, C. (2017). Truth, knowledge, and the standard of proof in criminal law. *Synthese*, <https://doi.org/10.1007/s11229-017-1608-4>.
- [17] McGrath, M. (2004). Review of John Hawthorne, *knowledge and lotteries*. *Notre Dame Philosophical Reviews*, vol. 8.

- [18] Moss, S. (2018a). *Probabilistic knowledge*. Oxford: Oxford University Press.
- [19] Moss, S. (2018b). Moral encroachment. *Proceedings of the Aristotelian Society*, 118(2), 177–205.
- [20] Moss, S. (forthcoming). Knowledge and legal proof. *Oxford Studies in Epistemology*, vol. 7. Oxford: Oxford University Press.
- [21] Nesson, C. (1979). Reasonable doubt and permissive inferences: the value of complexity. *Harvard Law Review*, 92(6), 1187–1225.
- [22] Nesson, C. (1985). The evidence or the event? On judicial proof and the acceptability of verdicts. *Harvard Law Review*, 98(7), 1357–1392.
- [23] Pardo, M. (2010). The Gettier problem and legal proof. *Legal Theory*, 16(1), 37–57.
- [24] Pardo, M. (2019). The paradoxes of legal proof: a critical guide. *Boston University Law Review*, 99(1), 233–290.
- [25] Pettit, P. (1989). Consequentialism and respect for persons. *Ethics*, 100(1), 116–126.
- [26] Pritchard, D. (2015). Risk. *Metaphilosophy*, 46(3), 436–461.
- [27] Redmayne, M. (2008). Exploring the proof paradoxes. *Legal Theory*, 14(4), 281–309.
- [28] Smartt, T. (2019). Review of Sarah Moss, *probabilistic knowledge*. *Ethics* 129(2), 430–438.
- [29] Smith, M. (2018). When does evidence suffice for conviction? *Mind*, 127(508), 1193–1218.
- [30] Thomson, J. J. (1986). Liability and individualized evidence. *Law and Contemporary Problems*, 49(3), 199–219.
- [31] Tribe, L. H. (1971). Trial by mathematics: precision and ritual in the legal process. *Harvard Law Review*, 84(6), 1329–1393.

- [32] Wasserman, D. T. (1991). The morality of statistical proof and the risk of mistaken liability. *Cardozo Law Review*, 13(2-3), 935–976.
- [33] Yablo, S. (2014). *Aboutness*. Princeton: Princeton University Press.