Episteme (2023), **20**, 367–376
doi:10.1017/epi.2021.45

ARTICLE

Statistical Evidence and the Problem of Specification

Frederick Schauer

University of Virginia, Charlottesville, VA, USA
Email: fschauer@law.virginia.edu

(Received 3 September 2020; revised 4 August 2021; accepted 14 September 2021;
first published online 22 November 2021)

Abstract

Philosophical debates over statistical evidence have long been framed and dominated by L. Jonathan Cohen's Paradox of the Gatecrasher and a related hypothetical example commonly called Prison Yard. These examples, however, raise an issue not discussed in the large and growing literature on statistical evidence – the question of what statistical evidence is supposed to be evidence of. In actual practice, the legal system does not start with a defendant and then attempt to determine if that defendant has committed some unspecified or under-specified act, as these examples appear to suppose. Rather, both criminal and civil litigation start with a sufficiently specified act and then attempt to determine if the defendant has committed it. And when we start with a more fully specified act, the statistics look very different, and these prominent examples no longer present the paradox they are claimed to support. Examining the issue of specification, however, does more than simply undercut the prominent examples in a long and extensive literature. The examination also raises normative issues challenging the legal system's traditional reluctance to base liability on the conjunction of probabilities.

Keywords: Statistical evidence; proof paradox; legal epistemology; specification; aggregate probabilities

1. Introduction – Cohen revisited

For more than forty years philosophers and academic lawyers have wrestled with the problem (or paradox) of statistical evidence: How is it that common intuitions and widespread legal practice resist finding defendants guilty (in criminal cases) or liable (in civil cases) on the basis of statistical evidence that appears to satisfy or surpass the applicable burden of proof? Why does the legal system allow judgments of guilt or liability based on imperfect eyewitness testimony, for example, but not allow judgments based on statistical evidence with even higher probabilities of accuracy?

Although foreshadowed by debates in the legal literature (Finkelstein and Fairley 1970, 1971; Tribe 1971), the *locus classicus* for philosophers has long been the so-called Paradox of the Gatecrasher offered by Cohen in *The Probable and the Provable* (Cohen 1977: 74–81). Almost every one of the literally hundreds of articles in legal and philosophical journals focuses on Cohen's seminal example, a hypothetical civil lawsuit that posits an event – a rodeo – in which 1000 people were in attendance as spectators, but in which only 499 of the spectators are determined to have paid for their admission. If

© The Author(s), 2021. Published by Cambridge University Press

https://doi.org/10.1017/epi.2021.45 Published online by Cambridge University Press

any one of the 1000 spectators were to be sued (civilly) for being a gatecrasher – for fraudulent entry – the probability that this spectator entered fraudulently is 501/1000, which exceeds the "preponderance of the evidence" (as it is put in the United States) or "balance of probabilities" (as often put in many British Commonwealth jurisdictions) burden of proof applicable in civil cases. It would seem, therefore, that this (and any other) randomly selected spectator should be held liable, but common intuitions and actual legal practice resist this conclusion. This tension between what the statistics say and what the legal system does creates the paradox. And although the occasional outlier has sought to defend the statistical conclusion (Shaviro 1989; Schauer 2003: 79–107), the vast bulk of the scholarship has attempted to demonstrate why it is correct that there is no legal liability in cases such as this, the seeming statistical conclusion notwithstanding (Thomson 1986; Moss 2018: 201–20; Di Bello 2019; Ross 2020).

Cohen's Paradox, however, is based on a misunderstanding of how the legal system actually operates, a misunderstanding that infects a raft of subsequent examples. Because this misunderstanding produces a mistaken statistical conclusion, correcting the misunderstanding in such an influential example can reorient for the better a debate whose interest among philosophers is, if anything, increasing (Smith 2018; Blome-Tillman 2020; Ross 2020). Correcting this misunderstanding, however, is only a secondary aim of this paper. The primary aim is to suggest that Cohen's (and others') misunderstanding of actual legal practice may open the door to a plausible but neglected normative possibility – the possibility that legal systems ought to be more willing than they now are to base liability on underspecified wrongs, at least when multiple underspecified acts are relevantly similar and are alleged to be committed by the same defendant.

2. Cohen's example – and others

Cohen's presentation of The Paradox of the Gatecrasher is brief enough to repeat in its entirety:

> Consider … a case in which it is common ground that 499 people paid for admission to a rodeo, and that 1,000 are counted on the seats, of whom A is one. Suppose no tickets were issued and there can be no testimony as to whether A paid for admission or [instead] climbed over the fence. So by any plausible criterion of mathematical probability there is a .501 probability … that he did not pay. The mathematicist theory would apparently imply that in such circumstances the rodeo organizers are entitled to judgement against A for the admission-money, since the balance of probability (and also the difference between prior and posterior probabilities) would lie in his favour. But it seems manifestly unjust that A should lose his case when there is an agreed mathematical probability of as high as .499 that he in fact paid for admission. (Cohen 1977: 75)

Cohen offers his example in the context of a civil lawsuit for monetary damages, but similar examples in the context of criminal prosecutions have been presented by others. Thus, Littlejohn (2017), presents a hypothetical case, Prison Yard, based on a similar hypothetical offered earlier by Nesson (1979) posing the same issue in the context of a criminal prosecution. As paraphrased by Backes (2020: 2760), the example goes as follows:

> 100 prisoners are exercising in a prison yard. Suddenly 99 of them attack the guard, putting into action a plan that the 100th prisoner knew nothing about.

https://doi.org/10.1017/epi.2021.45 Published online by Cambridge University Press

> The 100th prisoner played no role in the assault and could have done nothing to stop it. There is no further information that we can use to settle the question of any particular prisoner's involvement. … Suppose that after the event one of the 100 prisoners, chosen at random, was prosecuted for their involvement in the attack. Should this prisoner be convicted?

As with the Paradox of the Gatecrasher, Prison Yard is offered by many philosophers (Di Bello 2019) and legal theorists (Redmayne 2008) to demonstrate a mismatch between common intuitions and actual legal practice, on the one hand, and the result indicated by the statistics, on the other. And with few exceptions, the philosophical literature leans heavily in the direction of defending the intuitions against the statistics.

3. Evidence of what?

The aim of this paper is decidedly not to attempt to resolve all of the proof paradoxes and the longstanding debates they have spawned. It is, however, to argue that at least one paradox – the paradox seemingly created by both of these examples, and by many similar ones – is based on a misunderstanding of how the legal system operates. Simply put, both of the above examples proceed from the premise that we start with some defendant (whether randomly selected or not) and then attempt to determine whether that defendant has committed some liability-generating act. But the premise is descriptively mistaken. When Cohen's rodeo organizer sues the randomly selected spectator, the organizer must show not that this entrant is more likely than not – $p > 0.5$ – to have *somehow* entered fraudulently. Rather, the organizer must show that the entrant has committed a *particular* and specified fraudulent entry. And because there were 501 fraudulent entries, the fact that a defendant was in attendance shows only, absent other evidence, that the defendant was 1/1000 likely to have committed a particular liability-generating act. The probability that the defendant has committed *this* act is thus, without other evidence, not 0.501, but 0.001 – plainly insufficient to ground liability.

The same analysis applies to the Prison Yard scenario. As with Cohen's original example, the paradox as commonly presented embeds a descriptively false picture of legal procedure. Thus, Prison Yard presupposes that people can be convicted of a crime when there is a high probability that they have committed at least one of some number of unspecified offenses, even if we do not know exactly which one. Although there are plausible normative arguments for the legal system to operate in this way, as I will argue below, in fact the legal system does not now so operate. And because common law legal systems require that defendants be charged with particular crimes, and not merely with one or more unspecified tokens of a type of crime, the Prison Yard example never gets off the ground.

In Prison Yard, what we can label the *problem of specification* arises because the example asks us to imagine a prisoner picked at random and then tried. But what is the prisoner being tried for? Or, to frame the issue precisely, what is it that the prisoner is 99% likely to have done? The quick answer is that the prisoner is 99% likely to have been a participant in the murder of the guard. But how? Is he being charged with having delivered one out of 99 blows? Or inflicting one out of 99 knife wounds? Or standing watch while others actually struck the guard? Or instructing the participants as to each one's particular role? Or something else?

"So what?," we might ask. But "So what?" is the wrong question to ask in legal proceedings as currently structured, whether in a criminal case charging assault or a civil case like Cohen's charging some variety of fraud. In order to be found legally

https://doi.org/10.1017/epi.2021.45 Published online by Cambridge University Press

responsible, whether civilly or criminally, the randomly selected defendant must be charged with having engaged in a particular and specified unlawful act. If we do not know which act, or if we cannot describe and specify the act the defendant is alleged to have committed, he cannot be tried at all.[1]

If the prosecutor in Prison Yard were actually able to charge a particular act, then the probabilities look very different. If a randomly selected prisoner is charged with, say, knifing the guard in the chest, then it is not enough for the prosecutor to prove that there were 99 prisoners somehow involved in an assault. Assuming that there was only one knife wound, then the chances that the defendant committed *this* knifing are now 1/100. And the same holds true if there were multiple knife wounds, because it would still be incumbent on the prosecutor first to specify a particular knifing, and then attempt to prove that the defendant committed it. The 99% figure that allegedly creates the paradox is based on the false assumption that we start with the defendant and then seek to determine whether he committed some unspecified illegal act. But this is backwards. In reality, the legal system starts with a specified illegal act, and *then* attempts to prove that the defendant committed it. And as the difference between 0.01 and 0.99 shows, it makes all the difference whether we begin with the act or begin with the defendant.

Thus, what Cohen with Gatecrasher and those who have used Prison Yard see as a paradox is generated by an artifact of the example – the constructed fortuity of a large number of distinct acts and distinct bases for liability being highly similar. But similar or not, they are still distinct acts. Cohen's Gatecrasher example is especially seductive because all of the 501 fraudulent entries are highly similar, and the similarity produces the intuition that gets Cohen's paradox going. The similarity obscures the fact that each fraudulent entry is different, each occurring at least at a different moment. And because each fraudulent entry occupies a different point in space and time, we can see that Cohen hypothesizes his randomly selected defendant with having been sued for having engaged in an unspecified fraudulent entry, which the legal system as it exists does not permit.

To make the point clearer, consider the following, similar to an example in Pundik (2015), and which we can label Burglar: There have been ten burglaries in the city over the previous six months. Sam has recently been spending extravagantly, and in the presence of five witnesses at a bar he is asked where he suddenly obtained all that money. He responds that he broke into a house and stole a large quantity of cash and jewels. When asked if he was responsible for all ten recent burglaries he responds, "No, just one. I didn't want to push my luck." A witness reports this conversation to the police, who question Sam, but Sam refuses to provide any more information. Assuming no evidence other than Sam's confession to an unspecified act and his recent spending, can he be convicted of burglary? Can he be convicted of having committed one of the burglaries

[1]The relevant legal doctrine is described in Capra and Richter (2018) and defended in, inter alia, Clermont (2012) and Pundik (2015) A prominent exception is the California Supreme Court's decision in *Sindell v. Abbott Laboratories* (1980), allowing imposition of market-share liability in a products liability claim with multiple defendants, each of which marketed some of the allegedly defective product. But the fact that *Sindell* is controversial (Fennell 2018) reinforces the prevailing view, to which *Sindell* is the exception. Under the prevailing view, even products liability plaintiffs must show so-called *specific causation*, and merely showing (probabilistic) general causation is insufficient (Schauer and Spellman 2020). And although in criminal cases, some legal systems (the United States, notably) allow collective culpability by means of prosecutions for conspiracy and the like, even those prosecutions require a showing of precisely what a defendant is alleged to have done, and that is so even when what defendants have done make them legally responsible for the acts of others.

https://doi.org/10.1017/epi.2021.45 Published online by Cambridge University Press

that took place, even without proof as to which one? And here the intuition – no – is consistent with how the law operates.

The Burglar example differs from Cohen's Gatecrasher (and Prison Yard) only as a matter of degree. And once we recognize that each of the fraudulent entries in Cohen's Gatecrasher paradox is a separate event in the same way that each of burglaries in Burglar is a separate event, the intuitions go in opposite directions. And that is precisely because it becomes clear that the probabilities are nowhere near the high ones that Cohen posits, making his seeming paradox evaporate.

4. A different paradox still persists

As I suggested above, pointing out the necessity of specification does not solve all of the proof paradoxes and doing so is not my aim here. Still, because Cohen's Paradox of the Gatecrasher has framed and dominated a vast literature for more than forty years, with variants of the Prison Yard scenario being almost as important for almost as long, exposing the flaw in the example has value in its own right. More specifically, the Paradox of the Gatecrasher appears to be a paradox only because of a seeming intuition that there is something amiss with imposing liability on a randomly selected entrant for having committed an unspecified or under-specified wrong. But once we see that the legal system does not impose liability for such a spectacularly underspecified wrong, we can see that the intuition that leads Cohen and others to think there was a paradox is an intuition misled by an artifact of the example. As Burglar shows, once the artificial similarity of different wrongs is removed, the intuitions go in a different direction, and what Cohen supposes to be paradox is dissolved.

But perhaps Cohen and his worry should not be saddled with Cohen's misconceived example. Consider, for example, the almost-as-famous problem of the Blue Bus (Tribe 1971) and the real Massachusetts Supreme Judicial Court decision on which it was based (*Smith v. Rapid Transit Co.* 1945). Here we have the liability-producing act of a particular bus causing a particular accident in a particular way, but we do not know who owned the bus. Now the wrong appears therefore to be more than adequately specified, but the question remains as to whether a particular bus company can be held liable solely on the basis of statistical evidence that it operated most (well over 0.51) of the buses operating on this route at this time. Here the statistics incline towards liability but common intuitions still incline against. And these intuitions against liability cannot be so easily written off as generated by under-specification of the wrong. But *these* intuitions present a paradox different from that in Gatecrasher, even if no less a paradox. Here the paradox is presented not by an under-specified liability-generating act, but, rather, by an under-specified description of the defendant, and this is a different issue from that of under-specification of the wrong. The literature is mistaken in taking Gatecrasher and Prison Yard to present a problem, as just argued, but Blue Bus illustrates that a different problem of statistical evidence exists even when the wrong is adequately specified but the description of the defendant is not. Here the intuition that appears to create the problem is a function of under-specification of a quite different variety, and the paradox presented by *that* form of under-specification is not one I address here.

Putting aside, therefore, the potentially genuine problem of under-specification of the defendant, recognizing the under-specification of the liability-generating act shows Gatecrasher and Prison Yard to be less paradoxical than is widely assumed. But the recognition that appears to dissolve these and similar paradoxes exposes normative questions about just why the legal system and widespread intuition

https://doi.org/10.1017/epi.2021.45 Published online by Cambridge University Press

actually do demand precise specification of a liability-producing act. And to that I now turn.

5. Is act-specification necessary?

The argument presented has been based on what legal systems actually do, and, albeit to a much lesser extent, on digging deeper into the intuitions that people actually have. And with the occasional exception, the requirement of act-specification as I have formulated it is a common feature of both actual practice and widespread intuition. Exposing the issue of act-specification, however, invites consideration of whether such a requirement is justifiable.

Consider, therefore, a different example, which we can call Predator: John has been accused by four different women of sexual assault. Each accuser is credible, but so too is John's emphatic denial in each case. And for each accusation there are neither witnesses (other than the accuser and the accused) nor corroborating evidence. Under such circumstances, we might hypothesize a likelihood of John's guilt in each case as something in the order of 0.80. And then, assuming that proof beyond a reasonable doubt translates into something greater than 0.90 (Gardiner 2017), John will be properly acquitted in each of four independent prosecutions.

But now let us ask a different question: What is the probability, on these facts, and assuming independence among the chargers and the charges, that John has committed at least one sexual assault? Now the calculation is $1 - ((1 - 0.80) \times (1 - 0.80) \times (1 - 0.80) \times (1 - 0.80))$, which is 99.984. On these probabilities, therefore, there is a 99.984% chance that John has committed at least one sexual assault, even though there is only an 80% chance that he has committed a particular specified sexual assault, and even though he will be properly acquitted in each of four separate prosecutions. The question is whether John should be convicted of having committed a sexual assault.

The point of the example is to suggest that it is hardly obvious that the legal system is correct in requiring prosecution only of individually and precisely specified offenses (Harel and Porat 2009). Indeed, convicting John under these circumstances does not appear to run a risk (except a minuscule one) of sanctioning someone who is actually innocent, in much the same way that sanctioning Sam in the burglary case appears to involve no risk at all of sanctioning someone actually innocent.[2] By contrast, prosecuting everyone at the scene in the Prison Yard case guarantees that one innocent person will be sanctioned even though the evidence of guilt is very strong in each individual case. And much the same applies to Gatecrasher.

So then the central question is before us. As a normative matter, why is the legal system reluctant to follow the statistics and prosecute in cases such as John's? There is no question that each of the acts charged is sufficient to ground liability if proved. And thus, as a statistical matter, it is overwhelmingly likely that John has committed a liability-generating act, even if we do not know exactly which one. Indeed, it is overwhelmingly likely that John has committed a particular kind of liability-generating act, thus making the case for liability seem even more powerful.

A different example may make the point even stronger. Suppose the speed limit on a particular limited-access highway between Allenville and Brookstown, which are 130 miles apart, is 65 miles per hour. And suppose that Susan, as recorded by an automatic camera or at an electronic toll booth, enters the highway at Allenville at 9 a.m. and exits that highway on the same day at Brookstown at 10:45 a.m. In order to travel the 130 miles in 105 minutes, as Susan did, she must have averaged 1.24 miles per minute,

[2]Thanks to Sarah Moss for raising this point with me.

https://doi.org/10.1017/epi.2021.45 Published online by Cambridge University Press

which is slightly over 74 miles per hour. So Susan must have been speeding at some number of points. Perhaps she was driving at a steady 74, in which case she violated the law at every instant of her journey. But perhaps instead she traveled the first 65 miles in one hour, and thus entirely in obedience of the speed limit, but traveled the last 65 miles in 45 minutes, which is 1.44 miles per minute, or slightly over 86 miles per hour, this violating the speed limit for these 45 minutes. But we do not know where or when Susan was speeding, and thus we cannot specify the time and place when Susan violated the law. As with John the Predator, we are highly confident (actually, for Susan, certain) that Susan broke the law, but the precise act which violated the law is under-specified.

In the case of Susan, I suspect that a common intuition, and putting aside any concern with this method of electronic law enforcement generally, is that prosecution would be entirely permissible even in the absence of specification of the exact times and places of the speeding. But now let us change the example slightly. Suppose that Allenville and Brookstown are 100 miles apart, that the speed limit on the non-urban road connecting them is 50 miles per hour, and that there are five "Stop" signs between the two cities. And suppose that Susan is recorded at having left Allenville at 10:00 a.m. and arrived in Brookstown two hours later. Susan could have done this by exceeding the speed limit for some part of her journey but still stopping at all of the "Stop" signs, but she could have done this by driving at a steady 50 miles per hour and illegally stopping at none of the "Stop" signs. But she could not have done this without breaking the law in some way. So can Susan be prosecuted for having either exceeded the speed limit or for not stopping at one or more "Stop" signs, but without the prosecutor specifying which illegal act Susan has committed?

Now I suspect that the intuitions are likely somewhat different. And if the intuitions for the second Susan case are different from those for the first Susan case, the difference can only be explained by the fact that in the second case we are dealing with different types of unspecified offenses and in the first with unspecified tokens of the same type of offense. Indeed, if we return to the case of John the Predator, we can vary that case to make the point clearer. Suppose that instead of being charged by four different women with four different cases of sexual assault, each with a probability less than beyond a reasonable doubt, John is now charged with four different offenses. One person charges John with sexual assault, another charges him with embezzlement, a third accuses him of simple non-sexual assault, and a fourth alleges that John has committed a burglary. Each of the charges, again, is proved to a 0.80 probability. So now, again, there is a 0.9924 probability that John has committed a serious offense, but none are proved beyond a reasonable doubt. Can John be sanctioned for having committed a crime?

The sense that it would be wrong to sanction John here leads to the even more obvious *reductio* that if John can be prosecuted here, then pretty much any of us could be sanctioned based on the probability that we have committed some crime in the past, even if the authorities cannot specify which one.

If it is right to punish Susan in the first Susan case and perhaps also John in Predator, but not right to punish Susan in the second Susan case, not right to punish Sam in Burglar, and not right to punish John for having committed an unspecified instance of four very different types of crime, the conclusion is that it is right to require a specification of the type of wrong for which liability is sought, but not as obviously right that specification of precisely which token of a type of liability-generating wrong is necessary to produce liability.

Implicit in existing legal practice is the premise that the legal system should punish for acts and not punish actors except for committing particular acts. But the examples offered here expose that existing practice may depend upon a particular kind of act-

https://doi.org/10.1017/epi.2021.45 Published online by Cambridge University Press

specification, or, as sometimes put, *individuation* (Davidson 1969). Some wrongs, for example, are cumulations of multiple individual acts, or cumulations of patterns of acts. The wrong of creating a hostile environment in the workplace would be one example, and so too with the operating of an ongoing criminal organization, made an independent crime by the Racketeer Influenced and Corrupt Organizations Act (RICO) of 1970. More mundanely, hitting someone with a baseball bat nine times in immediate succession would be one criminal battery, but hitting that one person with the same bat once a day for nine consecutive days would be nine batteries – nine separate crimes.

6. On the legal metaphysics of act identity

The law itself provides no clear answer to the questions just posed, and that is largely because there is no clear answer. Still, if it can be an independent wrong to engage in the multiple acts that together constitute the single wrong of creating a hostile workplace environment, or if it can be an independent wrong to engage in the multiple chemical dumpings that together constitute the independent wrong of polluting a river, then it could, in theory, be an independent wrong to burgle the same house in more or less the same way on nine consecutive nights. And if this is so, then, and here the similarity with RICO is even closer, it could be an independent wrong to burgle nine different houses owned by nine different people on nine consecutive nights. And then it is no longer so clear that particularized specification is as essential as our intuitions suggest. There may be good reasons not to prosecute Sam in Burglar, or not to give Susan a citation in Susan II, but these would be legal and not epistemic reasons.

Fully examining and evaluating this issue would take us too far afield, in part because it would require exploring just why it is that we punish, and thus why it is that we punish acts and not actors. And at this point we are well-advised to consider an observation by Margalit, offered shortly after Cohen's book was published. Margalit argued that what Cohen perceived as a problem or a paradox was nothing of the kind, but merely a reflection of the fact that "there is more to proof in court than mere inductive reasoning" (Margalit 1980).

Margalit said no more than that on that occasion, but he was correct. There is more to legal proof than mere inductive reasoning,[3] and one component of Margalit's "more" is the longstanding requirement that a particular legal wrong be specified prior to determining whether some given agent has committed it. Once we see this, we can understand why Cohen's example has misfired. But we can also ask whether this long-standing requirement rests on sound normative foundations, a question that both statistics and the examples of John (Predator) and Susan may help us address. Moreover, the differences between Susan the speeder and Susan the unspecified traffic law violator illuminate the way in which the problem is not only one of statistics, and not only one of epistemology, but also one of what we might call legal metaphysics. So consider one final example, 7–11 Robber. Larry, wearing a ski mask and brandishing a pistol, approaches the clerk at a 7–11 convenience store and says, "Give me the money." The clerk then hands over the bills in the till one by one. Is this one robbery or many? And if the clerk hands over all the bills in the till and Larry then says, after a delay of several seconds, "Now give me the money in the safe," is this one robbery or two?

[3]Backes (2020) is correct (albeit with slight overstatement) that there "is no epistemic mileage in legal cases." And it is likely also true that there is little (even if not no) legal mileage in much of traditional epistemology (Enoch *et al.* 2012; Enoch & Fisher 2015).

https://doi.org/10.1017/epi.2021.45 Published online by Cambridge University Press

Implicit in law's traditional requirement of specification is the law's understanding of what makes a single act a single act, and when some combination of acts is a single act or instead multiple acts. Thomson (1977: 13) asks us to consider the example of "Sirhan killed Kennedy" as a way of arguing that questions of act-specification are often (or perhaps always) dependent on questions of causation, but the larger issue raised by Davidson, Thomson, and many others is about the way in which act-specification is domain- and purpose-dependent. The law has traditionally required a specification of the wrongs on which liability is premised, but probing just what counts as a single act for legal purposes makes it clear that even in law there is no bright line between single and multiple acts. And if it is right that Susan can be prosecuted for speeding without specification of exactly when and where she was speeding, and if people can be prosecuted for creating environmental hazards or hostile workplace environments, then perhaps it is not so clear that the law cannot and should not prosecute John in Predator.[4]

References

Backes M. (2020). 'Epistemology and the Law: Why There is no Epistemic Mileage in Legal Cases.' *Philosophical Studies* **177**, 2759–78.

Blome-Tillman M. (2020). 'Statistical Evidence, Normalcy, and the Gatecrasher Paradox.' *Mind* **129**, 563–77.

Capra D. and Richter L.L. (2018). 'Character Assassination: Amending Federal Rule of Evidence 404(B) to Protect Criminal Defendants.' *Columbia Law Review* **118**, 769–832.

Clermont K.M. (2012). 'Aggregation of Probabilities and Illogic.' *Georgia Law Review* **47**, 165–80.

Cohen L.J. (1977). *The Probable and the Provable*. Oxford: Oxford University Press.

Davidson D. (1969). 'The Individuation of Events.' In N. Rescher (ed.), *Essays in Honor of Carl G. Hempel*, pp. 216–34. Dordrecht: Reidel.

Di Bello M. (2019). 'Trial by Statistics: Is a High Probability of Guilt Enough to Convict?' *Mind* **128**, 1045–84.

Enoch D. and Fisher T. (2015). 'Sense and Sensitivity: Epistemic and Instrumental Approaches to Statistical Evidence.' *Stanford Law Review* **67**, 557–611.

Enoch D., Fisher T. and Spectre L. (2012). 'Statistical Evidence, Sensitivity, and the Legal Value of Knowledge.' *Philosophy & Public Affairs* **40**, 197–224.

Fennell L.A. (2018). 'Accidents and Aggregates.' *William & Mary Law Review* **59**, 2371–445.

Finkelstein M.O. and Fairley W.B. (1970). 'A Bayesian Approach to Identification Evidence.' *Harvard Law Review* **83**, 489–517.

Finkelstein M.O. and Fairley W.B. (1971). 'The Continuing Debate over Mathematics in the Law of Evidence.' *Harvard Law Review* **84**, 1801–9.

Gardiner G. (2017). 'In Defence of Reasonable Doubt.' *Journal of Applied Philosophy* **34**, 221–41.

Harel A. and Porat A. (2009). 'Aggregating Probabilities Across Cases: Criminal Responsibility for Unspecified Offenses.' *Minnesota Law Review* **94**, 261–310.

Littlejohn C. (2017). 'Truth, Knowledge, and the Standard of Proof in Criminal Law.' *Synthese*. http:doi.org/10.1007/s11229-017-1608-4.

Margalit A. (1980). 'Comments.' In L.J. Cohen and M. Hesse (eds), *Applications of Inductive Logic*, p. 195. Oxford: Clarendon Press.

Moss S. (2018). *Probabilistic Knowledge*. Oxford: Oxford University Press.

Nesson C. (1979). 'Reasonable Doubt and Permissive Inferences: The Value of Complexity.' *Harvard Law Review* **92**, 1187–225.

Pundik A. (2015). 'Should Criminals be Convicted of Unspecified Offenses? On Efficiency, Condemnation, and Cognitive Psychology.' *Criminal Law and Philosophy* **9**, 207–24.

[4]Earlier versions of this paper, or parts of it, have been presented at the World Congress on Evidential Reasoning in Gerona, the University of Oxford's Jurisprudence Discussion Group, the World Congress of the International Association for the Philosophy of Law and Social Philosophy, and the University of California at Los Angeles School of Law. I am grateful for audience comments on those occasions, as well as comments on drafts from Larry Alexander, Ruth Chang, David Enoch, Kim Ferzan, Valentin Hartmann, Sarah Moss, and Levi Spectre.

https://doi.org/10.1017/epi.2021.45 Published online by Cambridge University Press

Redmayne M. (2008). 'Exploring the Proof Paradoxes.' *Legal Theory* **14**, 281–309.
Ross L.D. (2020). 'Recent Work on the Proof Paradox.' *Philosophy Compass* **15**. doi.org/10.1111/phc3.12667.
Schauer F. (2003). *Profiles, Probabilities, and Stereotypes*. Cambridge, MA: Harvard University Press.
Schauer F. and Spellman B.A. (2020). 'Probabilistic Causation in the Law.' *Journal of Institutional and Theoretical Economics* **176**, 4–17.
Shaviro D. (1989). 'Statistical-probability Evidence and the Appearance of Justice.' *Harvard Law Review* **103**, 530–54.
Sindell v. Abbott Laboratories (1980). 26 Cal. 3d 588, 607 P.2d 924.
Smith M. (2018). 'When does evidence suffice to convict?' *Mind* **127**, 1193–218.
Smith v. Rapid Transit Co. (1945). 317 Mass, 469, 58 N.E.2d 754.
Thomson J.J. (1977). *Acts and Other Events*. Ithaca, NY: Cornell University Press.
Thomson J.J. (1986). 'Liability and Individualized Evidence.' *Law and Contemporary Problems* **49**, 199–219.
Tribe L. (1971). 'Trial by Mathematics: Precision and Ritual in the Legal Process.' *Harvard Law Review* **84**, 1329–93.

Frederick Schauer is David and Mary Harrison Distinguished Professor of Law at the University of Virginia. Formerly Frank Stanton Professor of the First Amendment at Harvard University and Professor of Law at the University of Michigan, his teaching and writing concentrates on evidence and proof, constitutional law, legal reasoning, the philosophy of law, and freedom of expression. His most recent book, *Proof: Uses of Evidence in Law, Politics, and Everything Else*, will be published by the Harvard University Press in the Spring of 2022.

Cite this article: Schauer F (2023). Statistical Evidence and the Problem of Specification. *Episteme* **20**, 367–376. https://doi.org/10.1017/epi.2021.45

https://doi.org/10.1017/epi.2021.45 Published online by Cambridge University Press