ORIGINAL ARTICLE

Justice in epistemic gaps: The 'proof paradox' revisited

Lewis Ross

London School of Economics and Political Science

Correspondence
Lewis Ross, Lakatos Building, Portugal Street, London School of Economics and Political Science, WC2A 2HE.
Email: l.ross2@lse.ac.uk

Abstract

This paper defends the heretical view that sometimes we ought to assign legal liability based on statistical evidence alone. Recent literature focuses on potential unfairness to the defending party if we rely on bare statistics. Here, I show that capitulating in response to 'epistemic gaps'—cases where there is a group of potential harmers but an absence of individuating evidence—can amount to a serious injustice against the party who has been harmed. Drawing on prominent civil law litigation involving pharmaceutical and industrial negligence, the overall aim is to illustrate moral pitfalls stemming from the popular idea that it is never appropriate to rely on bare statistics when settling a legal dispute.

1 | INTRODUCTION

A central question in applied legal philosophy is whether we should ever allow *bare statistics* to settle legal disputes. A swathe of recent work agrees that it is never appropriate to settle a legal case using statistical evidence alone or, alternatively, that it is only appropriate to do so when the odds are overwhelming such as DNA evidence with a < 1 in 10,000,000 chance of error.

Focusing on civil law, this paper defends the heretical opposing view: there are in fact cases in which relying on bare statistics is not only justifiable but demanded by justice, even if there is a substantial risk of error. I illustrate this argument by drawing on prominent civil law litigation involving pharmaceutical and industrial negligence in which bare statistics played a decisive role. A full appreciation of the cross-cutting considerations of justice that bear on legal verdicts should lead us not only to reject blanket criticism of bare statistics, but also to reconsider certain intuitions that underpin much of the debate—particularly those regarding the famous Blue Bus case, a case which has been taken to provide decisive support for the orthodox view for the better part of a century.

---

2 | THE AVERSION TO BARE STATISTICS

Recent philosophical work on legal statistics is motivated by a class of cases known as the 'proof paradox'. Proof paradoxical cases are scenarios in which the only evidence against the defending party (relative to some essential claim) is statistical.[1] Canonical examples, drawing on an older body of legal scholarship, are the following:

BLUE BUS: A bus negligently causes injury to a pedestrian, but it is not known which company the bus belongs to. On the route where the accident occurred, the *Blue Bus Company* runs 65% of the buses. There is no further information. [Adapted from Tribe, 1971]

PRISONERS: 100 prisoners are exercising in the prison yard. Suddenly 99 of them attack the guard, putting into action a plan that the 100th prisoner knew nothing about. 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 [Redmayne, 2008]

The BLUE BUS scenario concerns liability for a civil wrong—a negligent harm, typically called a 'tort'—while the PRISONERS case primarily concerns criminal wrongdoing, although many crimes can also be pursued as civil wrongs. In this paper the focus will be on civil law, although I will make some dialectical observations encompassing criminal cases.[2]

According to a popular view, it would be inappropriate to find against either the Blue Bus Company or any individual prisoner in these cases.[3] This common judgement is puzzling, especially if we focus on civil law. After all, the standard of proof operative in civil law is the *balance of probabilities*. This means that if the evidence supports a given fact on the balance of probabilities then that fact must be *treated as if it were true* by the court. The following judicial quote explains the idea well:

In ordinary (non-lawyers') language, to say that one regards something as 'probable' is by no means to say that one regards it as 'established' or 'proved'. Yet in the civil courts, where we say that a pursuer must prove his case on a balance of probabilities, what is held to be probable is treated as 'proved'. [Lord Prosser in *Dingley v The Chief Constable*, *Strathclyde Police* 1998 SC 548]

Since proof paradoxical cases seem to involve *by stipulation* evidence that supports liability on the balance of probabilities, how can we explain or justify our reluctance to rely on statistical evidence?

Many explanations have been proposed, often drawing from contemporary epistemology. For example, it has been claimed that bare statistics lack the right sort of *causal connection* to the truth (prominently Thomson 1986); that statistics do not rationally support *outright belief* (e.g. discussed by Buchak, 2014); that statistics fail to be *sensitive* to the truth in the right way (e.g. discussed by Enoch et al., 2012); that statistics are not *counterfactually safe* (e.g. see Pritchard, 2015; 2018 or Pardo 2018), that statistics do not provide *normic justification* (see Smith, 2018), that statistics fail to eliminate salient *error-possibilities* (see Gardiner 2019), or that statistics cannot generate legal *knowledge* (e.g. see Littlejohn, 2020 or Moss, 2016; forthcoming). These approaches claim there is a generic problem with bare statistics which can be traced back to some deficiency in its epistemic power. Almost all of these views are stated categorically, opposing reliance on bare statistics across the board, treating civil and criminal cases alike.[4] Some views make exceptions for statistical evidence involving truly miniscule chances of error—for instance, DNA evidence, which can be notionally accurate to a .9999999 degree of confidence—while upholding the rejection of bare statistics in familiar 'low odds' cases.[5] A defence of imposing liability in cases like BLUE BUS is heretical within recent philosophical discussion.[6]

A natural response is to return to a question posed by Enoch, Spectre and Fisher in their seminal (2012) paper. Paraphrased, the worry is: 'Fine, perhaps statistics lack certain epistemic properties. *But why should the law care*?' After all, following where bare statistical evidence leads is a reliable way to deliver the correct verdict, if we take a probabilistic interpretation of reliability. Indeed, bare statistics can be *more likely* to lead us to the truth than other types of evidence on which we routinely rely (consider the notorious unreliability of the eye-witness). The sharpest version of the challenge can be pressed as follows: why should the law countenance *accuracy sacrifices* by precluding verdicts based on bare statistics and valorising epistemic properties like knowledge?

One way to answer this question is to focus on the possibility that relying on bare statistics might inflict or risk inflicting an *injustice* upon the defending party. This strategy is particularly appealing when—as many authors do—we make criminal versions of the proof paradox the primary focus of the debate. For example, in the PRISONERS case, our mind is immediately drawn to the unhappy possibility of a prisoner suffering sanctions for a crime they are entirely innocent of. Particularly when criminal punishment is at stake, the failure to preclude these possibilities can strike us as deeply unfair—surely we need to do more in order rule out false convictions? Precisely this intuition has been seized upon in recent work. Pritchard (2018) suggests that bare statistics, being counterfactually unsafe, are morally problematic because they create too much of a certain type of *risk* of false convictions; Littlejohn (2020) suggests that only evidence which can yield legal *knowledge* satisfies moral principles concerning when blame can be considered reasonable; and Smith (2021) explores the idea that, in lieu of a certain type of epistemic *justification*, allowing bare statistics to win the day commits us to deliberately harming the innocent in certain scenarios.[7]

The current dialectic of the debate then is as follows: it has been widely noted that we can explain intuitive reluctance to rely on bare statistics by appealing to theories from epistemology that pinpoint the absence of some generic epistemic property. A natural way to explain why *this* epistemic property is so important for legal purposes is to suppose that it aligns with *moral* constraints on evidence law that require us to eschew bare statistics when deciding legal cases.

This analysis offers a satisfying way to connect epistemic and moral concerns. However, although this approach certainly directs the debate in the right direction, I contend that it is incomplete in various important respects: indeed, I claim that some of these respects are so important as to in fact recommend an entirely different solution to certain proof-paradoxical cases. The current literature has identified some important pro tanto moral concerns with basing a legal verdict on the basis of bare statistical evidence. However, by focusing only on injustice to the *defending* party, we fail to account for the existence of moral imperatives that speak against leaving a harmed party without redress. It is to these imperatives I now turn.

3 | THE INJUSTICE OF REJECTING STATISTICAL EVIDENCE

Doctrines of evidence law that prevent us from issuing a verdict on the basis of certain types of evidence are double-edged. Although precluding positive verdicts relying on some class of fallible evidence E removes the possibility that E can be used to undergird a *false* verdict, it also ensures that E cannot be used to secure a *correct* conviction or attribution of liability.

As our concern here is primarily with civil law, reconsider the BLUE BUS case. While it might seem unfair were the Blue Bus Company to be found responsible on the basis of mere statistical probability, the other side of the coin is that there is a party which has suffered injury and faced with the prospect of it going uncompensated. For, if all we have to go on are statistics, and we are

contemplating a rule on which bare statistics can never be sufficient to impose liability, then in such cases the claim of the *injured party* is doomed to failure. When a party is negligently harmed and cannot receive legal redress, this itself amounts to an injustice which goes unrectified. This observation should give us pause. If the *rejection* of bare statistics leads to an injustice of some description, the simple moral account offered earlier cannot be the full story. Hence, we need a broader conception of the sorts of considerations relevant to the debate on bare statistics. To begin, we will need to make some germane points about evidence law.

The policy-like ramifications of precluding verdicts based on certain types of evidence puts pressure on the idea that we can decide how to treat bare statistics simply by looking at generic facts about their epistemic power, without thinking about the broader implications of endorsing any such rule. What do I mean by the *policy-like* ramifications of evidence law? Let me explain. Although recent debates are motivated by individual 'proof-paradoxical' scenarios, in reality these cases would be decided with reference to the applicable tenets of evidence law. There are two aspects of evidence law worth noting here. Firstly, the relevant principles will be used not just to settle the case at hand but an entire class of relevantly similar cases. And secondly, evidence law can be altered (or left unaltered) in service of broader socio-juridical goals.[8] Putting these points together, we see that deciding upon the right principles of evidence law requires us to think about their broader effects multiplied over a number of similar cases. So, for example, we might decide to exclude (i.e. render inadmissible) evidence gained through police impropriety even though such evidence might be highly reliable in a given instance—however, the *cumulative* effect of admitting such evidence over many cases might serve to encourage police misconduct and hence incubate unfairness. As I have emphasised elsewhere, the debate on statistical evidence is not about admissibility: rather it is whether such evidence is *sufficient* to undergird a legal verdict. Nevertheless, if we decided to endorse a rule on which bare statistical evidence was deemed insufficient evidence across the board, then we would be pre-committing ourselves to favouring the defending party in *every single case* with structural similarities to BLUE BUS. The question then becomes whether the cumulative effect of such a rule might frustrate broader concerns of justice or social goals that we might expect the legal system to care about?

Of course, it is entirely appropriate to have a natural concern about revising *downwards* the epistemic standards that must be met before imposing a legal sanction. But this worry should not be overstated. Evidence law is malleable and evolves in response to perceived injustices within the trial process; there is nothing God-given about the current standards.[9] More importantly, adjusting downwards certain legal protections can be a way of dealing with *structural imbalances* surrounding the sufferance of certain harms. For instance, a very important debate concerns whether certain aspects of evidence law should be *weakened* in cases involving sexual assault as a way to address the perennially low conviction rates of such crimes. A guiding thought is that, due to engrained sexist attitudes, myths surrounding consent, and the uniquely private nature of many sexual offences, the law is in some sense 'stacked against' victims of sexual assault. Some argue that we should remove certain protections in favour of the defendant—for example, corroboration requirements in which two independent sources of evidence are needed for conviction[10]—to remedy this imbalance in sexual assault trials. For my purposes it is irrelevant where you sit on the first-order question about how exactly to deal with this aspect of criminal justice. We may, ultimately, reject such arguments about criminal practice. But, it is clearly legitimate to debate whether we ought to reconsider tenets of evidence law in order to address broader structural imbalances concerning how certain harms are processed in the legal system.

With these considerations about the double-edged nature of evidence law on the table, I will now turn to look more closely at the civil law before identifying two prominent common-law examples which illustrate that rejecting bare statistics can lead to serious injustice.

3.1 | A Closer Look at the Law

In keeping with the tenor of recent philosophical work, we have proceeded at a high level of generality. However, fully appreciating how legal developments relate to the philosophical debate surrounding statistical evidence will require a little more precision in discussing the nuts and bolts of the law.

Earlier we introduced the 'balance of probabilities' standard of proof operative in civil law. Moving beyond the standard gloss provided in philosophical discussions, it is important to now recognise that legal proof involves a number of different elements—*each* which must individually (rather than collectively) be established on the given standard of proof.[11]

To establish liability for a negligent harm, the pursuer must prove:

> **HARM**: The pursuer suffered *harm*.
> **DUTY OF CARE**: The defender had a *duty of* care towards the pursuing party.
> **NEGLIGENCE**: The defender acted *negligently* (breached their duty of care).
> **CAUSATION**: The defender's negligence *caused* the harm complained of.

In what follows, I focus primarily on CAUSATION. The conventional notion of causation incumbent on the pursuer to prove is a counterfactual notion that must be established through a simple 'but for' test—X is treated as causing Y if it is shown that Y would not have occurred but for X.[12]

Proving the causal element turns out to be an impossibility in certain important cases involving substantial negligent harms—cases in which we fall into what I term an *epistemic gap*. Bare statistics can play a crucial role as surrogate evidence in these cases. The cases I discuss below have structural similarities with BLUE BUS, without being entirely isomorphic. Nonetheless, considering them will furnish us with strong reasons to reconsider the currently dominant hostility to bare statistics. A secondary purpose of considering these cases is to highlight, contrary to a common refrain in the proof-paradox literature, that the law is no stranger to allowing bare statistics to carry the day in certain instances.

3.2 | Asbestos, Mesothelioma, and Heavy Industry Labourers

Asbestos is a naturally occurring mineral used sporadically by humans since the Stone Age but found intensive application as an insulator and fire retardant in the construction and shipbuilding industries. Evidence of the harmful nature of asbestos on the human respiratory system has existed since around 1900, but most jurisdictions only restricted asbestos near the end of that century. Asbestos exposure causes respiratory diseases such as mesothelioma which are painful, debilitating, and often fatal. Typical sufferers are working-class males once employed in heavy industry and exposed to asbestos dust in the workplace, and their partners who were exposed to asbestos dust on contaminated work clothes.

In civil cases brought by labourers attempting to win compensation for being negligently exposed to asbestos by their employers (e.g. due to a lack of protective equipment like face-masks), establishing a causal link between any particular period of asbestos exposure and contraction of illness was often impossible. Not only do we have an imperfect understanding of the underlying mechanisms generating diseases like mesothelioma, but crucially, a career in the heavy labour industry often involved transient employment by different corporations—in almost every case, victims had been exposed to asbestos over many years, by different employers, and in different contexts. Those bringing the cases were simply unable to adduce evidence establishing that any particular employer was responsible for causing the diseases that were subjecting them, in some instances, to a slow and lingering death. Evidence like: "A exposed B to asbestos dust for five years" does not satisfy the traditional 'but for' test that lawyers use to establish causation of a specific outcome such as a disease. (One might object that these cases involve joint causation. However, intriguingly, a popular medical theory at the time was that mesothelioma is not 'dose dependent' and can be caused by a single fibre of asbestos. In this sense, the assumption truly was that a single party caused the illness.)[13] There was an epistemic gap with respect to which particular employer's negligence caused the disease. As a result, the lawsuits brought by those who had been poisoned in this way were doomed to fail. Consequently, an entire class of people, these labourers being among the more vulnerable groups in society, were systematically unable to receive justice after contracting diseases while employed by large and extremely profitable multi-national corporations who demonstrated a marked indifference to their safety.

Epidemiological evidence—that is, *statistical evidence* about the risk of disease associated with exposure to asbestos dust[14]—was generally not considered to be a satisfactory way of establishing a *causal* connection between negligent exposure and eventual contraction of mesothelioma. However, in a series of cases beginning with *Fairchild v Glenhaven Funeral Services*[15], developed in *Barker v Corus*[16] and terminating in *Sienkiewicz* v *Greif*[17], the courts developed a series of principled exceptions to the usual requirements of legal proof in order to rectify this situation.[18] In short, the courts held that in asbestos cases where ordinary proof of causation is impossible, we can instead rely upon bare statistical evidence about the risk of disease resulting from a period of negligent exposure to asbestos. As a result, a negligent employer who had materially increased the risk of a given victim contracting mesothelioma could be held liable for causing the disease.

Specifically, the liability of the employers was 'joint and several': the pursuer could fully recover against any employer who had materially increased the risk of the eventual harm. Individual employers held liable, in turn, could recover from each other in proportion to the risk they had created. The advantage of joint and several liability is that it enables full redress even when some employers no longer exist as a legal entity.

The courts' reasoning here was plainly influenced by the thought that, as a matter of policy, it would be unjust to deny pursuers any legal recourse in these types of case. In the most influential of these cases, Lord Bingham concludes:

…there is a strong policy argument in favour of compensating those who have suffered grave harm, at the expense of their employers who owed them a duty to protect them against that very harm and failed to do so, when the harm can only have been caused by breach of that duty and when science does not permit the victim accurately to attribute, as between several employers, the precise responsibility for the harm he has suffered. I am of opinion that *such injustice as may be involved in imposing liability on a duty-breaking employer in these circumstances is heavily outweighed by the injustice of denying redress to a victim.*" Per Lord Bingham at para 33, *Fairchild v Glenhaven Funeral Services* [2002] UKHL 22

In endorsing such a position, the courts acknowledged that there might be reasons to worry about using bare statistics to impose liability—but held that these are outweighed by the importance of avoiding an approach that would systematically preclude industrial labourers from receiving compensation for asbestos-related illnesses.

3.3 | Carcinogenic Drugs and Market-share Liability

Throughout medical history there are instances in which a drug is developed and prescribed to alleviate some ailment, only for that drug to later be determined harmful in its own right. During the mid-20[th] century, a number of companies were involved in producing a synthetic oestrogen compound—diethylstilbesterol or 'DES'—that was given to pregnant woman as an anti-miscarriage drug. Around thirty years later it became public knowledge that DES was a cause of various cancers in the daughters of mothers who had taken the drug. It was further found that the companies involved in producing the drug had failed to test and advertise it in a responsible way, and therefore had been negligent in fulfilling their duty of care towards consumers.

As with the asbestos cases, proving the causal element of the tort was extremely difficult. The impediment was that a number of different companies had been involved in the manufacture of DES and it was not feasible, over thirty years later, for any individual sufferer to work out which of the various candidate companies had produced the particular tablets that had caused their cancer. Again, those who had been negligently harmed found themselves facing an epistemic gap with respect to causal responsibility. Therefore, the women bringing these cases were doomed to fail—to their extreme detriment as they would often be facing considerable healthcare costs, and to the benefit of the pharmaceutical corporations who had otherwise acted in a negligent way. Clearly, there was a public policy imperative that these women not be left without legal recourse.

The evidence that *was* available in these cases was a statistical analysis of the respective market-share of the companies involved in manufacturing the drug at the time the drug was being sold. In other words, market-share information provided the courts with a statistical probability that a given company was responsible for the harm caused to a given sufferer. In a very famous case—*Sindell v. Abbott Laboratories*[19]—this statistical evidence was used to develop the doctrine of '*market-share liability*'.[20] The essence of this doctrine was that liability could be assigned to the negligent companies on the basis of their respective market-share, even without there being any non-statistical evidence bearing on their causal responsibility. The result would be that these companies would *shoulder the compensatory burden in proportion to their share of the market* at the time the harm was caused. This was considered to be a fair outcome even though it guaranteed that, in any isolated case, a number of defenders would liable for harms which they did not cause. The upshot was that those afflicted with DES-related cancers could receive compensation.

3.4 | Summary

In these cases. the courts were faced with deciding whether to let bare statistical evidence carry the day. For ease of reference, we can reproduce simplified versions of these cases as follows:

DRUG: A pharmaceutical product that was negligently manufactured and advertised causes cancer in a consumer, but it is not known which out of a group of producers manufactured the particular pills taken by the consumer. In the area where the drug was purchased, the *Big Pharma Company* manufactured 65% of the relevant drugs on the market.

ASBESTOS: Fibres of asbestos that a labourer was negligently exposed to causes him to contract mesothelioma, but it is unknown which employer exposed him to the fibres which caused the disease. The *International Shipbuilding Company* was responsible for 65% of the labourer's exposure to asbestos fibres above the environmental norm, exacerbated by the fact that they, along with his other employers, failed to supply protective equipment.

Various courts judged that rather than acting fairly in disallowing bare statistics from establishing the essential causal element of legal proof, it would instead have been wrong to prohibit reliance on statistics.

There are important questions about these decisions. What is the theoretical basis for these judgements? And how far should we extend our willingness to attribute liability in this way? I discuss these issues in the next section. But before moving on, it bears emphasising that we have a very interesting result that goes against the tide of recent philosophical work on statistical evidence. If these legal doctrines are well-founded, then there should be no *categorical* bar against deciding an essential claim on the basis of bare statistics simply in virtue of unfairness to *defending* parties. And relatedly, nor should there be any universal bar on relying on bare statistics simply in virtue of them lacking some epistemic property, even when the chance of error is relatively high. Rather, if the cases described were defensibly decided, we have uncovered important instances in which bare statistical evidence can legitimately underpin legal liability.

4 | JUSTICE AND EPISTEMIC GAPS

I have outlined cases in which the pursuer was confronted with an *epistemic gap* that prevented them from establishing causation of harm. These situations have the following schematic features:
Epistemic Gaps

 (i) it is granted that harm has been wrongfully inflicted;
 (ii) we can identify a privileged reference-class of potential harmers;
(iii) it is impossible for the pursuer to individuate liability using familiar types of evidence that confer knowledge.

Such epistemic gaps can occur in both criminal and civil law. Sometimes epistemic gaps are entirely irremediable, but in the ASBESTOS and DRUG cases there was another form of evidence—namely statistical evidence concerning epidemiology and market-share—available to fill the epistemic gap and attribute liability. The question is whether the court was justified in filling the epistemic gap with recourse to statistical evidence?

Lord Bingham framed the decision to rely on statistics in terms of *justice*. I think that this is essentially correct, but what notion of justice might vindicate and explain this conclusion? Contemporary philosophy is overwhelmingly concerned with *distributive* justice. However, the idea that distributive concerns can properly justify doctrines of tort law is controversial among legal theorists. Indeed, a common idea is that tort law is essentially concerned with a second type of justice, namely *corrective* justice.

Corrective justice has a long pedigree. For example, alongside distributive justice it was the second category of justice identified by Aristotle.[21] Roughly, corrective justice concerns the restoration of the notional parity that exists between parties prior to their entering into some relationship. The criteria against which corrective justice is measured is entirely to do with how things were, and how things could have been, between the relevant parties to a relationship. For

example, if I culpably fail to keep to some contract between myself and a third party, the criteria for restitution—that is the criteria for *correcting* the wrong—is simply compensating the third party for whatever loss is due to my contractual failure. This can be contrasted with compensating them according to some independent distributive standard. Clearly then, corrective and distributive justice can pull in different directions. For example, if Tiny Tim dishonestly accrues some financial benefit at the expense of a wealthy tycoon then this would be something that a concern with corrective justice would enjoin us to rectify, even if Tiny Tim's negligent benefit resulted in a more just *distribution* of resources when judged against some plausible standard of distributive justice.

I now argue that attributing liability in the ASBESTOS and DRUG cases is vindicated by *both* distributive and corrective conceptions of justice, in addition to having a strong forward-looking rationale.

Corrective Justice

Corrective justice, taking the position of the respective parties prior to their relationship as the benchmark, concerns correcting harms wrongfully imposed during that relationship. If a given company *C* had negligently caused (for instance) the cancer of some party to whom they owe a duty of care (such as a consumer), then this would uncontroversially constitute a wrong within the ambit of corrective justice. In the cases under discussion, there is uncertainty about *which* negligent party caused the harm.

I suggest that using statistics to impose liability effectively *approximates* corrective justice. This is easiest to see with the DRUG case and the doctrine of market-share liability. The relevant question is: how does relying on market-share statistics leave us with respect to corrective justice over the *class* of cases to which the doctrine applies? Rejecting market-share liability, obviously, ensures a perfect absence of corrective justice insofar as no negligently harmed party will receive recompense. By contrast, apportioning liability according to market-share should, *over the class of cases*, approximate corrective justice. The victims will receive recompense, and pharmaceutical companies collectively be held responsible in proportion to their contribution to the overall harm. So long as we have no reason to suppose that this mechanism will disfavour any particular drug company—e.g. because consumers of one company have some higher propensity to sue than others—then liability by market-share is an effective way of ensuring that each defending party will be on the hook only for roughly their share of the overall harm caused. It is worth noting that the use of statistics in such cases is *defeasible*—defending parties are free to adduce evidence that there are specific reasons why market-share statistics are a misleading guide to how much harm they are responsible for.

This 'approximation' argument depends on some weak assumptions about the relative disvalue of two types of corrective *in*justice:

1. A pursuer's inability to receive compensation for a harm.
2. A defender erroneously being held liable for a harm.

Approaches like market-share liability all but prospectively guarantee instances of (2). After all, the relationship between market-share and harm caused is not a perfect correlation; these drugs did not have a 100% rate of causing cancer. We accept this imperfection as a way of eliminating *many more* instances of (1). Hence, justifying market-share approaches by appealing to

the approximation of corrective justice depends on the supposition that the disvalue of (2) is not orders of magnitude greater than (1). It isn't feasible to mount a full-scale defence of this supposition here, but we can note that this supposition is contained within the logic of civil law. Recall, the relevant standard on which the pursuer must prove causation is the *balance of probabilities* rather than some higher standard. This suggests a tacit commitment to something in the neighbourhood of parity of value between false positives and false negatives.[22] (Contrast, for illustration, the criminal law, where the 'beyond a reasonable doubt' standard clearly embodies the idea that false positives—i.e. false convictions—are considerably worse than false negatives.) Hence, there is a plausible argument from the approximation of corrective justice to rely on bare statistics in establishing causation in some legal contexts.[23]

Distributive Justice

While appealing to distributive concerns in advocating for positions in civil law is controversial, there are good arguments that such concerns must play some role. It is certainly not a category error to criticise aspects of civil law on the grounds that it is regressive with respect to distributive justice (e.g. see Keren-Paz, 2018). Moreover, it has been persuasively argued that we can legitimately appeal to distributive concerns in deciding *which* corrective injustices are to be given legal recognition (e.g. Gardner, 2014).

I there is a fairly compelling line of thought on which distributive concerns are relevant to epistemic gap cases. The distributive implications of an institutional structure are typically an extremely important metric of the overall justice of any institution. More specifically, in scenarios like ASBESTOS and DRUG, the court is facing circumstances in which both courses of action—either letting the epistemic gap sink the case *or* relying on 'non-causal' statistical evidence—are non-ideal. From the perspective of the court, this is a *forced-choice* scenario. There is no ideal judicial solution that avoids both types of concern. Given the existence of such non-ideal cases, where perfect corrective justice may not be possible, it seems reasonable for a court to at least consider distributive concerns in order to adjudicate on which of two non-ideal courses of action should be taken. In sum, I suggest that there is a prima facie case for the relevance of considerations of distributive justice to epistemic gap cases; the burden of proof must rest with opponents of this view to establish the irrelevance of distributive justice.[24]

If distributive concerns are properly relevant, then the decision to attribute liability in the ASBESTOS and DRUG cases can be vindicated by a concern for distributive justice. Take the mesothelioma cases as an example. The court faced a choice between **(A)** a rule that shields shipbuilding corporations from liability for their negligence and **(B)** a rule that enables heavy industry labourers suffering from negligently inflicted mesothelioma to be compensated for these harms. The latter rule was chosen. The distributive argument—as with the change to the law of evidence mooted in ASBESTOS and DRUG—range over case *classes* rather than just individual cases. It seems clear that the solution adopted by the courts will cumulatively tend towards a more just distribution of resources on any plausible standard of distributive justice. In general, distributive justice will be served by directing resources to the class of peripatetic labourers at the expense of multinational corporations, where the court is faced with an either/or choice between these outcomes.

We have now attempted to vindicate the use of statistics on both grounds of distributive and corrective justice. However, before moving on, it is worth flagging a third family of considerations relevant to justifying legal doctrines. These are *forward-looking* considerations that concern the effect of a legal rule on society generally.

Forward-looking Considerations

Forward-looking considerations often motivate judicial decision-making. Many forward-looking considerations broadly concern the *incentives* that the existence of a given legal rule would create for those potentially affected by it. While the proper role of such considerations is controversial, it is clearly compelling to suppose that the creation of beneficial incentives is a good-making feature of a legal rule. For instance, a tort that holds a shopkeeper liable for merchandise falling on top of customers incentivises shopkeepers to take precautions; this is uncontroversially a benefit of having such a rule. Indeed, the relevance of incentives is accepted by opponents of statistical evidence who have constructed arguments claiming that reliance on bare statistics would create perverse incentives in certain situations (most influentially, see Enoch et al., 2012).

With this in mind, we can state the obvious: using bare statistical evidence to attribute liability in the ASBESTOS and DRUG cases provides employers with an incentive to improve safety standards. For example, in some asbestos cases, the negligence of the employer consisted in their refusal to provide showers to labourers caked in asbestos dust. If it were known that there could be no liability for such negligence due to the lack of causal evidence, there would be no legal incentive for them to remedy such negligence. The reliance on statistical evidence in cases like ASBESTOS ensures that employers have reason to raise their safety standards. Indeed, the possibility of irremediable epistemic gaps could create *perverse* incentives to avoid taking certain precautions in the knowledge that the subsequent harms could not be causally attributed to any particular entity. The importance of appropriate incentives provides an argument in favour of the recourse to statistical evidence in filling some epistemic gaps.

A final forward-looking consideration concerns the effect of cases like ASBESTOS and DRUG on the public perception of justice in the legal system.[25] It has been argued that the mere perception of injustice is reason to be reluctant to rely on bare statistics in certain criminal cases. In the cases we are discussing, this consideration militates in the opposite direction. If the courts were forced to throw out every asbestos-related or DES case, there would be a public perception of *in*justice. To the extent that the public perception of the legal system has any role in justifying our approach to bare statistics, it seems to favour their use here.

To summarise: the use of bare statistics to settle a legal case does not, across the board, constitute an injustice against the defending party. To the contrary, we have identified prominent examples from the civil law—justified on distributive, corrective, and forward-looking grounds—where bare statistics carried the day. Crucially, to refuse to rely on bare statistics in these cases would have constituted an injustice against the *pursuing* party.

5 | BLUE BUS: AN ORTHODOXY REVISITED

Now that we have made some progress in understanding why eschewing bare statistics can create injustice, I want to return to the BLUE BUS case. It has become canonical to suppose that relying on bare statistics in BLUE BUS is woefully misguided. However, there are important similarities between it and the DRUG and ASBESTOS cases. So, is the received view really as robust as typically assumed?

I want to begin by suggesting that our intuitions against relying on statistics in cases similar to BLUE BUS are not categorical. Consider the following scenario:

MONOPOLY BUS: A bus negligently causes injury to a pedestrian, but there was no eye-witness evidence linking the bus to a particular company. On the route where the accident occurred, only one outfit has a regular service: the *Monopoly Bus Company*. Uncontested statistics from analysing CCTV cameras in adjacent neighbourhoods show that only 1 in every 10,000 buses passing through that area is owned by a private individual. There is no further information.

On standard critical accounts of bare statistics, the evidence against the *Monopoly Bus Company* is just as impotent as the evidence in the regular version of the BLUE BUS case: it lacks all of the epistemic properties outlined earlier (safety, sensitivity, normalcy, etc.). But failing to discern *any* relevant difference between these cases, I think, is incorrect.[26] Cases involving near-complete monopolies throw into sharp relief the fact that refusing to attribute liability on the basis of bare statistics is at odds with the ideals of distributive and corrective justice.[27] In this sense, there is a stronger justice-based argument for attributing liability in MONOPOLY BUS than in BLUE BUS. However, the difference between these cases is a matter of degree rather than kind. As such, I think that the question of whether to attribute liability in proof-paradoxcial scenarios depends on weighing the relevant reasons rather than entertaining any categorical prohibition on statistical evidence.

I will now argue that the strength of the reasons favouring attributing liability in BLUE BUS have been underappreciated. For ease of discussion, let's focus on a version where there are only two bus companies in the relevant area:

BLUE BUS: A bus negligently causes injury to a pedestrian, but it is not known which company the bus belongs to. On the route where the accident occurred, the *Blue Bus Company* runs 65% of the buses and the *Red Bus Company* runs 35% of the buses. There is no further information.

We are left with a choice between two sub-optimal options—impose liability without causal evidence or leave injured parties in such cases uncompensated. This is an important similarity between the BLUE BUS, ASBESTOS and DRUG cases. However, there is a crucial structural difference in the nature of the epistemic gap found in BLUE BUS compared to the ASBESTOS and DRUG cases. In the latter, *every* potential causer of harm—i.e. each heavy industry employer, and each drug manufacturer—had been negligent. By contrast, in the BLUE BUS case, only one member of the reference-class was negligent. The uncertainty in BLUE BUS concerns *both* which of the relevant parties was negligent *and* which of the relevant parties caused the harm. In the ASBESTOS and DRUG cases there was no risk of sanctioning a non-negligent party, only a causally inefficacious one. In BLUE BUS, if we used some mechanism to favour the pursuer, we risk sanctioning a non-negligent company. This structural difference raises once more the spectre of unfairness to defending parties.

I will now suggest that legitimate concerns about the unfairness of imposing liability in BLUE BUS are outweighed by countervailing considerations in favour of imposing liability. For sake of discussion, I will focus on the possibility of apportioning liability according to market-share.

Starting with considerations relating to corrective justice, it seems evident that attributing liability according to market-share gets us closer to the ideal of corrective justice than throwing out the case. This will obviously be true if we assume that there will be iterations of the BLUE BUS scenario with the two companies being causally efficacious in different instances, but bearing the costs for each accident according to their market-share. Over repeated iterations, market-share liability, all else being equal, will tend towards a perfect approximation of corrective justice. However, we will also get closer to corrective justice even if we assume that the BLUE BUS case is a one-off. Legally, the magnitude of a corrective injustice is quantified in monetary terms according to the level of damages that would be owed to the pursuing party. Suppose that this is assessed at *n*. Throwing out the case for want of evidence guarantees that the remaining corrective injustice

is of magnitude n. Attributing liability according to market-share, regardless of which company in fact causes the harm, will lead to a corrective injustice of a smaller magnitude than n. (If the Blue Bus Company causes the accident then there is a corrective injustice against the Red Bus Company of $n*.35$ and a corrective injustice of $n*.65$ if the Red Bus Company caused the crash. Each lower than the corrective injustice of n that occurs if the case is thrown out).

Moving to distributive justice, it is reasonable to suppose that commercial transport firms will be better able to bear unexpected financial losses in proportion to their market-share, or to bear relevant insurance costs, than individual pedestrians. Of course, as with the distributive justification for the decision in the ASBESTOS and DRUG cases, this must be interpreted as pertaining to the *class* of cases rather than necessarily to each individual case. Interpreted this way, the distributive case for market-share liability in BLUE BUS is straightforward.

There may still be a lingering sense of iniquity in using market-share liability in cases involving competing parties where not all have been legally negligent. In cases like DRUG, all of the potential causers of harm—i.e. all of the DES manufacturers—were just as bad as each other. They had each been negligent, so we feel little compunction about holding each of them market-share liable. This common normative shortcoming appears to be absent in BLUE BUS-type cases, where only one party has been negligent. I suggest that we can assuage this worry by considering some attractive supplementary normative principles.

Consider the following principle:

SHARED STANDARDS: Whenever a harm is negligently caused but falls into an epistemic gap, it is reasonable to apportion responsibility among potential harmers when they each share similar *standards* when conducting the risky activity.

The initial worry with attributing liability in BLUE BUS was that the parties were not equally responsible for a normative shortcoming, i.e. actually having caused harm in a negligent fashion. But, market-share liability can justifiably be predicated on the defending parties having roughly similar *standards* with respect to the risk of causing negligent harm. Hence, in such cases, there can nonetheless be an important sense in which there is normative parity between the defending parties. When there is such rough parity, we can justify apportioning liability in the BLUE BUS case by noting that it roughly reflects the amount of *risk of negligent harm* that a given party is exposing pedestrians to.[28] The important thing to remember about market-share liability is its defeasibility. Those 'on the hook' in virtue of their market-share can adduce evidence in their favour demonstrating that they have more exacting safety standards, blocking the inference from $n\%$ share of the market to $n\%$ likelihood of having caused the harm.

Focusing on standards in this way has important advantages. Firstly, it has the forward-looking benefit of creating a virtuous circle for companies worried about such liability to increase their standards above those of their immediate competitors. Secondly, and more fundamentally, it reduces the influence of *luck* in determining which party assumes liability. The problematic role of luck in tort law is a familiar concern—the very same negligent action can have vastly different results depending on circumstantial luck. Consider the following counterfactual. A bus driver negligently pulls out without looking. World A: no-one is nearby, no damage is caused. World B: she dents a Ferrari, causing £50,000 of damage. Circumstantial luck can generate extreme variation in the sanctions a party is liable for. By apportioning liability according to market-share in the BLUE BUS case—in other words, apportioning liability according to the level of risk imposed—we reduce the influence of such circumstantial luck.

A second principle further supports the legitimacy of attributing liability in BLUE BUS:

RESPONSIBILITY FOR EPISTEMIC GAPS: when parties can reasonably be held responsible for the *existence* of an epistemic gap, this weighs in favour of them bearing (part) responsibility for harms falling into that epistemic gap.

Modern technology affords companies with ways to gather extremely reliable evidence about where their vehicles are at any given time. Trackers are used for a variety of purposes by companies such as to provide additional information to consumers (e.g. about expected delivery times) and to monitor the performance of their employees. By the same token, such trackers could establish the involvement or otherwise of a given vehicle in an accident. Whether or not to install these devices, and to accept the consequences of their not doing so, is a decision entirely up to commercial agents. Given that this option exists, and that it would not be particularly onerous for the commercial outfits to take it, their decision not to close the epistemic gap can, I suggest, justify favouring the pursuer in a forced-choice situation where the pursuer is harmed by the very existence of the epistemic gap. Such a principle has the obvious forward-looking benefit, where applicable, of incentivising parties to take steps to eliminate epistemic gaps.

Of course, as has been pointed out in the literature on liability to defensive harming, even innocent victims of harm play some role in determining how much risk of harm there is in a given scenario.[29] An individual could eliminate risk by increasing their own safety standards in virtue of remaining in bed all day. Equally, they could eliminate any epistemic gap by carrying a GPS tracker at all times. My argument here, as with many in the law, turns on a notion of reasonableness—who is it reasonable to expect to take steps to increase standards and eliminate epistemic gaps? In all the cases discussed, I submit, that burden lies with the defending party. These defeasible principles—**SHARED STANDARDS** and **RESPONSIBILITY FOR EPISTEMIC GAPS**—militate in favour of attributing liability by market-share in the venerable BLUE BUS case.

To close, I want to consider an objection to my argument concerning the relevance of alternative compensation mechanisms. The orthodox response to the proof paradox enjoins us to treat the harms in BLUE BUS just like 'Acts of God'—such as lightning strikes—for which there is no legal remedy. Some might argue that there should ideally be state-provisioned compensation for victims of Blue Bus-style epistemic gaps that ameliorate the harsh effect of rejecting such cases. It is far from clear to me that a taxpayer-funded compensation scheme would be the normatively ideal response to such cases since it would involve the public assuming responsibility for corporate negligence. Perhaps other models, such as an industry-financed compensation scheme, would be more defensible. However, there is a sense in which this is irrelevant to the normative debate concerning how the courts should adjudicate on epistemic gap cases when they are faced with them. For, it is not within the gift of the courts to institute such a scheme. The debate on legal statistics has become increasingly concerned with making prescriptions about how courts should act in the face of real legal dilemmas—as it borne out by recent claims about the proper use of DNA evidence—so the positions we adopt in this debate must respond to the nature of the society in which we live. In the ASBESTOS and DRUG cases, the courts were faced with a situation in which other compensation mechanisms did not exist. My argument has been that in attributing liability, they made the right decision. Of course, it may be that defenders of orthodox solutions to the proof paradox deny this and hold that the courts should have rejected such cases. Perhaps they believe that compensating people for harms that fall within epistemic gaps is a legislative responsibility and never a judicial one. But, if this is so, then it turns out their solution to the proof paradox rests on a fundamental jurisprudential position about the proper role of the courts which has not yet been defended as an assumption in this debate. This, by itself, is a theoretically striking finding.[30]

The BLUE BUS case has been taken as a paradigmatic demonstration of why we should never rely on bare statistics. I hope to have shown that things are far less clear than is often supposed. Not only does it have important affinities with other legal cases in which attributing liability is entirely defensible, but there are viable arguments—from distributive justice, corrective justice, and forward-looking considerations—supporting the use of statistics about market-share to attribute liability in the BLUE BUS case.

6 | CONCLUDING REMARKS

Justice can demand that the law rely on bare statistics in attributing liability, as shown by prominent legal responses to gaps in evidence surrounding asbestos poisoning and carcinogenic medicines. Not only does this go against an emerging orthodoxy in legal philosophy about the solution to the proof paradox, but it also underscores broader lessons for the debate on statistical evidence and legal philosophy generally.

Firstly, while it is natural to focus on how using suboptimal evidence can wrong a defending party, there are powerful cases in which a failure to accept 'second best' evidence constitutes an injustice against those who have been harmed. The law of evidence and proof are double-edged; rejecting certain types of evidence may make it harder to get things wrong, but it also makes it harder to get things right too. Especially in the civil law, when there can be vast inequalities in the resources available to pursuers and defenders, there are situations in which justice demands that we err on the side of accepting false positives rather than false negatives.

Secondly, future research must go beyond the currently dominant approach to statistical evidence on which universal prescriptions are made about the proper role of statistics based on considering a narrow range of cases. In particular we must go beyond justifying such universal prescriptions by noting *generic facts* about the epistemic properties of such evidence. This approach leaves out the rich context in which rules of evidence evolve, the differences between criminal and civil law, and the fact that different cases bring up bespoke issues. Moreover, cases involving bare statistic are structurally various—sometimes they involve a collection of negligent parties, sometimes they involve a single negligent party, and there are likely unexplored cases between these extremes.

Thirdly, doctrines of evidence law are amenable to discussion from the lens of distributive justice, a fact that recent work on statistical evidence has not accounted for. Not only does this further support the contention that we cannot make a priori universal prescriptions about the proper role of bare statistics, it provides us with a way to make important distinctions between different cases. Put bluntly, there are different imperatives of justice when large and wealthy organisations do battle in court compared to the case of a moribund sufferer of asbestos-related mesothelioma attempting to gain compensation from their employer.

Finally, it is worth stressing that I do not mean to suggest that contemporary work on bare statistical evidence reaches the wrong conclusion in the majority of cases. But it *is* important to now move towards a theoretical stance that moves beyond categorical prescriptions and displays sensitivity to the different considerations relevant to different types of case. Not only will this bring legal nuance into the philosophical literature, it will enable us to contribute more effectively to socially important dilemmas faced by the legal system. In this way, I hope, we can continue the rich vein of legal theory produced by philosophers in recent years.[31]

ENDNOTES

[1] For simplicity I often use the Scots' law terms 'pursuer' and 'defender' when referring to the sides in a legal case. Pursuer is equivalent to 'plaintiff' (in US law) or 'claimant' (in English law).

[2] The important differences between criminal and civil law are many. Firstly, false positives are widely considered less acceptable in the criminal context. Secondly, it is harder to reconcile the use of bare statistics with the criminal 'beyond a reasonable doubt' standard of proof than with the civil 'balance of probabilities' standard. And thirdly, while an important argument against bare statistics in the criminal law has turned on supposing that probabilistic evidence cannot sanction beliefs or assertions involving moral blame, it is controversial whether moral blame plays any role in civil judgments. My own view is that the standards of criminal proof have a different justificatory basis than the standards of civil proof, one requiring the former to be interpreted in a particularly exacting way (Ross ms).

[3] Not only philosophers have these judgements; they have been empirically vindicated by psychologists under the guise of the 'Wells' effect', so named due to Wells (1992).

[4] There is also a legal literature which attempts to vindicate the aversion to bare statistics by offering new ways to understand the standard of proof (e.g. Cheng, 2013) and the nature of legal fact-finding (e.g. Sullivan, 2019). While this approach is predominant, there has been some dissent. See Egglestone (1980) for an early rejection of the puzzle. Hedden and Colyvan (2019) defend a probabilistic conception of the standard of proof against the idea that bare statistics are particularly paradoxical; Ross (2021a) argues against the supposed parallel between individualistic and legal epistemology; Krauss (2020) offers a legally-informed rebuttal of recent criticism of bare statistics.

[5] For example, see Cheng and Nunn (2016), Enoch and Fisher (2019), Di Bello (2019).

[6] See Ross (2021b) for an alternative criticism of this orthodoxy, appealing to cases involving conjunctions of different types of statistical evidence.

[7] Legal scholars have also appealed to moral strategies (or something like them) to argue against bare statistics: see Wasserman (1992), Stein (2005), and more recently Nunn (2015). See Redmayne (2008: 292-296) and Pundik (2008) for discussion.

[8] See Stein (2005) for a lawyer's perspective on these issues.

[9] E.g. see Laudan (2006) for an argument that the 'beyond reasonable doubt' standard in criminal law is too favourable to defendants and fails to exhibit enough concern for the danger posed by repeat offenders. Gardiner (2016) provides statistically-informed criticism. I also reject Laudan's argument in Ross (ms).

[10] The requirement of corroboration is found most prominently in Scots' law. While the corroboration rule has been retained in general, exceptions have been introduced that relax this requirement.

[11] Readers might notice that the *overall probability* of each of these being true might fall <.5 even if each is established to a >.5 likelihood. This is known as the 'conjunction paradox' and working out how best to react to it is a perennial puzzle in legal theory. See Pardo (2019: 266-282) for a survey of possible responses.

[12] As with philosophy of causation, there are hard cases that have prompted alternative proposals of legal causation. See Moore (2019) for survey and critique.

[13] The 'single fibre' theory is now less popular.

[14] Epidemiological evidence is structurally similar to statistical evidence in other proof-paradoxical cases, involving a reference-class (humans exposed to asbestos dust) and a frequency (instances of disease in members of that reference-class). The difference is that it involves multiple reference-classes: (i) exposure to merely environmental levels of asbestos, and (ii) exposure to varying occupational levels of asbestos. By comparing the two, it is possible to get an estimated increase in risk of disease as a result of occupational exposure.

[15] [2002] UKHL 22

[16] [2006] 2 A.C. 572

[17] [2011] UKSC 10

[18] For legal commentary, see Steel and Ibbetson (2011). I should flag that legal interpretation of these cases is controversial. For my purposes, these controversies are not especially important: I submit that my reading is normatively plausible irrespective of disagreements about how to interpret the judicial dicta.

[19] 26 Cal. 3d 588

[20] A paragraph in Krauss (2020) also highlights the doctrine of market-share liability as bearing on the proof-paradox. An older paper by a legal scholar, mostly overlooked by philosophers, also flags the similarity between

DES cases and the proof-paradox. In that paper, the author expresses scepticism about whether "the depth of the pockets we assume to be present" should have any bearing on the correct result (Brooks 1980: 344-5). Below, I argue that he is incorrect.

[21] See *Nicomachean Ethics*, Book V. For Aristotle, distributive justice concerned divisible things like honours and goods, while corrective injustice concerned voluntary and involuntary transactions. See Weinrib (2012) for contemporary discussion; Sheinman (2014) for discussion of the relationship between corrective and distributive justice in tort law.

[22] It is possible to argue that civil law exhibits a slight preference for false negatives over false positives, since multiple claims have to be established on the balance of probabilities for a civil case to succeed. There is considerable lack of clarity on this issue from a theoretical perspective, illustrated by debates on the 'conjunction paradox'.

[23] Admittedly, there may well be contextual aspects relevant for comparing the disvalue of (1) and (2) that vary with the details of the case.

[24] Before concluding the paper, I end by considering an objection from the role of the legislature in providing compensation.

[25] See Tribe (1971); Brooks (1980); Ross (2021a) for further discussion of statistics and public perception.

[26] Empirical study supports my intuition. Subsequent examination of the Wells' effect found it to be largely absent in cases involving 99.9% probabilities (see Wright et. al 1996).

[27] Of course, as with any trial, the defending party would be free to adduce exculpatory evidence that undermines the relevance of the statistics—we are not seeking to impose strict liability onto monopolies.

[28] I flag, without endorsing, a separate idea according to which risk-imposition can amount to harm in its own right. See Oberdiek (2017) for discussion.

[29] See, for example, Lazar (2009) for an extensive discussion.

[30] See Stein (2005) for useful discussion.

[31] My thanks to Marcello Di Bello and Ho Hock Lai for comments, to audiences at the University of Glasgow's *Law and Philosophy Network* and at the London School of Economics' *Popper Seminar* for sharp and generous questioning, to my friends and colleagues in the LSE's Probe Group for their feedback on a previous draft, and to everyone else with whom I have discussed the proof paradox over the past few years.

REFERENCES

Blome-Tillmann, Michael (2015). Sensitivity, Causality, and Statistical Evidence in Courts of Law. *Thought*, *4*, 102–112. https://doi.org/10.1002/tht3.163

——— (2017). 'More Likely Than Not' - Knowledge First and the Role of Statistical Evidence in Courts of Law. In Adam Carter, Emma Gordon & Benjamin Jarvis (eds.), *Knowledge First - Approaches in Epistemology and Mind*. Oxford: Oxford University Press. pp. 278–292.

Brook, J. (1984). The use of statistical evidence of identification in civil litigation: Well-worn hypotheticals, real cases, and controversy. *Saint Louis University Law Journal*, *29*, 293–352.

Buchak, Lara (2014). Belief, credence, and norms. *Philosophical Studies*, *169*, 285–311. https://doi.org/10.1007/s11098-013-0182-y

Cheng, Edward (2013). Reconceptualising the Burden of Proof. *Yale Law Journal*, *122*, 1254–1279.

——— & Nunn, G.A. (2016). DNA, Blue Bus, and phase changes. *The International Journal of Evidence & Proof*, *20*, 12–120. https://doi.org/10.1177/1365712715623556.

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

Di Bello, Marcello (2019). Trial by Statistics: Is a High Probability of Guilt Enough to Convict? *Mind 128*, 1045–1084. https://doi.org/10.1093/mind/fzy026

Eggleston, Richard (1980). The Probability Debate. *Criminal Law Review*, *Nov*, 678–688.

Enoch, David, Spectre, Levi, & Fisher, Talia (2012). Statistical Evidence, Sensitivity, and the Legal Value of Knowledge. *Philosophy and Public Affairs*, *40*, 197–224. https://doi.org/10.1111/papa.12000

Gardiner, Georgi (2017). In defence of reasonable doubt. *Journal of Applied Philosophy*, *34*, 221–241. https://doi.org/10.1111/japp.12173

——— (2018). Legal Burdens of Proof and Statistical Evidence. In David Coady & James Chase (eds.), *The Routledge Handbook of Applied Epistemology*. London: Routledge.

――― (2019). The Reasonable and the Relevant: Legal Standards of Proof. *Philosophy & Public Affairs*, *47*, 288–318. https://doi.org/10.1111/papa.12149

Gardner, J. (2014). What is tort law for? Part 2. The place of distributive justice. In Oberdiek, J. (Ed.). *Philosophical foundations of the law of torts*. Oxford University Press.

Hedden, Brian & Colyvan, Mark (2019). Legal Probabilism: A Qualified Defence. *Journal of Political Philosophy*, *27*, 448–468. https://doi.org/10.1111/jopp.12180

Krauss, Sam (2020). Against the Alleged Insufficiency of Statistical Evidence. *Florida State University Law Review*, *47*, 801–825.

Keren-Paz, T. (2018). *Torts, egalitarianism and distributive justice*. Routledge.

Laudan, Larry (2006). *Truth, Error, and Criminal Law: An Essay in Legal Epistemology*. Cambridge: Cambridge University Press.

Lazar, Seth (2009). Responsibility, Risk, and Killing in Self-Defense. *Ethics*, *119*, 699–728. https://doi.org/10.1086/605727

Littlejohn, Clayton (2020). Truth, Knowledge, and the Standard of Proof in Criminal Law. *Synthese*, *197*, 5253–5286. https://doi.org/10.1007/s11229-017-1608-4

Oberdiek, J. (ed.) (2014). *Philosophical foundations of the law of torts*. Oxford University Press.

――― (2017). *Imposing Risk: A Normative Framework*. Oxford: Oxford University Press.

Pritchard, Duncan (2015). Risk. *Metaphilosophy*, *46*, 436–461. https://doi.org/10.1111/meta.12142

――― *(*2018). Legal risk, legal evidence and the arithmetic of criminal justice. *Jurisprudence*, *46*, 436–461. https://doi.org/10.1080/20403313.2017.1352323

Moore, Michael, "Causation in the Law", *The Stanford Encyclopedia of Philosophy* (Winter 2019 Edition), Edward N. Zalta (ed.)

Moss, Sarah (2016). *Probabilistic Knowledge*. Oxford University Press.

――― (2021). Knowledge and Legal Proof. In *Oxford Studies in Epistemology Vol. 7*. Oxford: Oxford University Press.

Nunn, Alex (2015). The Incompatibility of Due Process and Naked Statistical Evidence. *Vanderbilt Law Review*, *68*, 1407–1433.

Pardo, Michael (2019). The Paradoxes of Legal Proof: A Critical Guide. *Boston University Law Review*, *99*, 233–290.

Pundik, Amit (2008). Statistical Evidence and Individual Litigants: A Reconsideration of Wasserman's Argument from Autonomy. *The International Journal of Evidence & Proof*, *12*, 303–324. https://doi.org/10.1350/ijep.2008.12.4.304

Redmayne, Mike (2008). Exploring the Proof Paradoxes. *Legal Theory*, *14*, 281–309. https://doi.org/10.1017/S1352325208080117

Roth, Andrea (2010). Safety in Numbers - Deciding When DNA Alone is Enough to Convict, *NYU Law Review*, *85* 1130–1185.

Ross, Lewis (2020). Recent Work on the Proof Paradox. *Philosophy Compass*, *15*, https://doi.org/10.1111/phc3.12667

――― (2021a). Rehabilitating Statistical Evidence. *Philosophy and Phenomenological Research*, *102*, 3–23. https://doi.org/10.1111/phpr.12622

――― (2021b). Legal proof and statistical conjunctions. *Philosophical Studies*, *178*, 2021–2041. https://doi.org/10.1007/s11098-020-01521-z

――― *(ms) Criminal Proof: Fixed or Flexible? See* http://lewisdylanross.com/papers/

Sheinman, H. (2014) Tort Law and Distributive Justice. In Oberdiek, J. (Ed.). *Philosophical foundations of the law of torts*. Oxford University Press.

Steel, S., & Ibbetson, D. (2011). More grief on uncertain causation in tort. *The Cambridge Law Journal*, *70*, 451–468. https://doi.org/10.1017/S0008197311000572

Stein, Alex (2005). *Foundations of evidence law*. Oxford: Oxford University Press.

Smith, Martin (2018). When Does Evidence Suffice for Conviction? *Mind*, *127*, 1193–1218. https://doi.org/10.1093/mind/fzx026.

――― (2021). Against Legal Probabilism. In J. Robson and Z. Hoskins (eds). *Truth and Trial*. Routledge. https://doi.org/10.4324/9780429283123

Sullivan, Sean (2019). A Likelihood Story: The Theory of Legal Fact-Finding. *University of Colorado Law Review*, *90*, 1–66.

Thomson, Judith Jarvis (1986). Liability and Individualized Evidence. *Law and Contemporary Problems*, *49*, 199–219. https://doi.org/10.2307/1191633

Tribe, L. H. (1971). Trial by mathematics: Precision and ritual in the legal process. *Harvard Law Review*, *84*, 1329–1393.

Wasserman, D. T. (1992). The Morality of Statistical Proof and the Risk of Mistaken Liability. *Cardozo Law Review*, *13*, 935–976.

Weinrib, Ernest Joseph (2012). *Corrective Justice*. Oxford: Oxford University Press.

Wells, Gary (1992). Naked statistical evidence of liability: Is subjective probability enough? *Journal of Personality and Social Psychology*, *62*, 739–752. https://doi.org/10.1037/0022-3514.62.5.739

Wright, Edward, Lora Maceachern, Elaine Stoffer, & Nancy Macdonald (1996). Factors Affecting the Use of Naked Statistical Evidence of Liability, *The Journal of Social Psychology*, *136*, 677–685. https://doi.org/10.1080/00224545.1996.9712244