

Statistical Evidence and Incentives in the Law

Is it bad practice to use statistical evidence as the basis for a finding of guilt or liability? And if it is bad practice, why is it bad practice?

This paper focuses on the incentivizing aspects of candidate legal systems in civil cases. Incentives clearly matter—a legal system which encouraged bad behavior could easily be bad, even if it reliably penalized bad behavior and did not penalize good behavior. And the use of statistical evidence does have implications for incentives.

We are by no means the first authors to explore the connections between statistical evidence and incentives. In two much cited papers, Enoch, Fisher, and Spectre (2012) and Enoch and Fisher (2015) consider this issue, contending that it is bad practice to use statistical evidence because the use of statistical evidence incentivises bad behavior.¹ They connect their claims about incentives to claims about epistemology—that the perverse incentives go hand-in-hand with the relevant legal findings being insensitive, that is, failing key counterfactual tests on tracking the truth.

As we said, thinking about incentives when evaluating the merits of using statistical evidence is certainly a good idea, and we think there is a good deal to be learned by careful engagement with these two papers. But, as we explain in Section 1, Enoch et al. go wrong in their evaluations of incentives. After introducing some important distinctions, we explain in some detail what incentives there are in various classic cases from the statistical evidence literature. In Section 2, we examine the relationship between incentives and epistemological sensitivity. We argue that the picture provided in the Enoch papers is misleading, and that legal incentives have no special connection to epistemological sensitivity. Section 3 discusses an important class of cases – toxic torts – and explain why they provide a much better incentive based case for findings of liability based on statistical evidence.

¹ The incentive approach is also taken more recently by Christian Dahlmann in 'Naked Statistical Evidence and Evidence for Lawful Conduct' (2020). We shall also make some observations about that paper in our discussion. The literature contains a profusion of approaches to the question of what's wrong with convicting on the basis of statistical evidence, offering a wide range of moral, epistemological, practical and economic considerations against the practice. (See, for example, Cohen 1977, Nesson 1979, Thomson 1986, Wasserman 1991, Blome-Tillmann 2015, Nunn 2015, Posner 1999, Smith 2017, DiBello 2018, Gardiner 2018, Moss 2018, Pardo 2019.) We shall note a few places where that work interfaces with an incentive based approach. But it is beyond the scope of this paper to critically engage with alternative approaches.

1. Sanchirico on predictive evidence and trace evidence.

Enoch and al.'s theory of incentives draws significantly from Sanchirico's work on character evidence (where character evidence is about an individual's propensities for engaging in criminal behavior). As a prelude to looking at that theory, we wish first to draw attention to some important aspects of Sanchirico's discussion—not all of which are represented correctly in Enoch et al.'s discussion.

1. Predictive evidence and trace evidence.

The distinction between predictive evidence and trace evidence plays a central role in Sanchirico's discussion. That distinction is inspired by H. Richard Uviller (1982), which contrasts trace methods and predictive methods for determining whether a certain event occurred at some time. The trace method consists of looking for traces of the event that are causally downstream from the event. The predictive method consists of looking for facts that obtained before the time and that were predictors of the event. Here is Uviller:

To illustrate: the trace method of proving that a nail was struck by a hammer is to discover the head of the nail and there discover the trace of a hammer blow. Another way is to establish that a certain person, being a carpenter, has frequently hit nails with his hammer and, having had the opportunity in question, and motive to do so, he probably struck the nail in question. (p. 847)

Drawing on contrasts like this, Sanchirico distinguishes trace evidence and predictive evidence. Predictive evidence

describes the setting – broadly defined to include props character, scenery – in which the conduct may or may not have occurred. Certain circumstances, including not only the actual physical situation but also the character or disposition of the parties involved, may be “fertile ground” for conduct; other settings may be particularly infertile.²

For our purposes, the following gloss will suffice: The predictive evidence possessed by a group regarding certain alleged conduct consists of facts that are known to obtain prior (sometimes just prior) to the alleged conduct and which have some probabilistic bearing on the likelihood of that conduct having been performed or not performed.³ By contrast, trace evidence has to do with what is causally downstream from the conduct:

² Sanchirico (2001), p. 1260.

³ The term 'evidence' is sometimes applied to facts, sometimes to objects and events. Both uses are acceptable in ordinary language. For our theoretical purposes it is better to focus on the fact construction as the theoretical apparatus of probabilities is better suited to propositions than to objects and events.

Trace evidence is evidence that is generated (or, more realistically, tends to be generated) by the conduct in question. Such evidence includes the consistent and “cross-resistant” testimony of multiple eyewitnesses of commission or its likely aftermath, authenticable documentation, fingerprints, bank account balances, DNA evidence, etc.⁴

Sanchirico’s elaboration—“or more realistically, tends to be generated”—is important. Suppose someone reports having seen Jones burn down a barn. In fact the eyewitness is lying, trying to frame Jones. As it happens the barn caught fire due to an electrical fault. The eyewitness report is not causally downstream from Jones’ burning the barn down, since Jones did not burn the barn down. But Sanchirico wants to count the eyewitness report as trace evidence because it’s the sort of that that would “tend to be generated” by a fact corresponding to the report. We think it slightly cleaner not to finesse matters this way: We shall say that a fact is trace evidence to a group for an act at *t* just in case the fact is the causal upshot of (or causally explained by) the act and which bears probabilistically on the likelihood of the act. (In principle, trace evidence can lower the probability that the act was performed. But in all the cases of interest it will raise the probability.) We can then, *inter alia*, recognize another category – a fact that is not in fact trace evidence, but which bears probabilistically on whether the act was performed due to the probability that the fact is a causal upshot of the act. Let us use ‘putative trace evidence’ as the umbrella term for trace evidence plus this latter category.

Note that the distinction between trace evidence and predictive evidence should not be confused with the distinction that is often made between individualized evidence and non-individualized evidence, which we can intuitively gloss as evidence that is *about* the individual that is the suspect of some crime and evidence that is not.⁵ Whether some evidence is predictive or trace does not settle whether it is individualized. Suppose Jones burns down his neighbor’s barn. Among the predictive evidence might be that Jones hated his neighbor (individualized) and that barns have in the past burned down more often because of arson than because of electrical faults (non-individualized). Among the trace evidence might be that an eyewitness claimed to have seen Jones burning down the barn (individualized) and that the charred remains of a barn were on the ground the next morning (non-individualized).⁶

⁴ Sanchirico (2001), p. 1261.

⁵ Thomson 1986 suggests that “what people feel the lack of and call individualized evidence is evidence that is in an appropriate way causally connected with the (putative fact) that the defendant caused the harm” (203). It is hard to think that this is quite right. In the case where Jones burns down a barn, the charred remains are surely causally connected with the putative fact that Jones burned down the barn. But we do not anticipate anyone calling the fact that there are charred remains of a barn “individualized evidence”.

⁶ Some complications remain regarding Sanchirico’s distinction between trace evidence and predictive evidence. Suppose *x* is accused of a violent act against *y* on Tuesday. On Wednesday, *z*, upon hearing about the alleged crime, makes a phone call that reports on violent acts that *x* had committed towards them on days prior Tuesday. The phone call occurs

2. Incentives

Sanchirico argues that the issue of whether and how to allow character evidence is best approached by widening one's gaze from the truth-seeking objectives of trials. True enough, a key goal of the legal system is that we don't punish the innocent in criminal trials, and we don't make people or institutions pay compensation for damages they were not responsible for in civil cases. Thus there is certainly a desire for accuracy in legal proceedings, and rules of evidence can certainly be evaluated according to how accuracy-conducive they are. But, as Sanchirico emphasizes, the legal system also aims to regulate the activity of citizens, and rules of evidence and can also be evaluated in terms of their influence on those activities by looking at the incentive structures they create.

We can illustrate the contrast using a fantastical but instructive case discussed in Sanchirico's paper (we add a few details and change a few others, but the essential structure is a version of the second "weather" case that he describes).⁷ Suppose there is an interest in people keeping their umbrellas down when it is sunny. To this end, a law is made criminalizing the putting up of umbrellas when the sky is blue. Suppose further that the legal system has excellent access to whether the sky was blue on a certain day, but disdains trace evidence altogether. Instead it adopts a very eccentric protocol: The society knows of a generic association between the presence of a certain chemical in the brain and a tendency to derive pleasure from putting up umbrellas when it is sunny. They have scanners that determine whether a person has that chemical in their brain. Let's suppose each citizen gets scanned once a month, and is told the result of the scan right away. If during the month the sky is blue, people who scanned positively for the chemical receive a fine in the mail.

There are certain things obviously wrong with this legal system. Scanning on a regular basis is likely to be intrusive and expensive. But the point that is relevant for our purposes is that the scanning system creates a hopeless incentive structure. As Sanchirico emphasizes, no one has

after Tuesday, and it is causally downstream of the crime (since hearing of the crime induces the phone call), but it seems awkward to classify it as trace evidence. What we learn here is that some evidence that is causally downstream of a crime may be of evidential value by being evidence of the kinds of propositions that are predictive evidence -- their evidential value would be screened off if one were already certain of the relevant predictive evidence. Suppose meanwhile that the phone call has no causal connection to the violent act. It is not then trace evidence. But nor is it predictive evidence since it occurred after that fact. The distinction between what is causally downstream from a fact and what is not is not the same as the distinction between what is causally downstream from an act and what pre-dates it. The latter distinction is arguably more fundamental for legal purposes than the predictive/trace distinction, though the previous point suggests that the causally downstream/non-causally downstream is not quite adequate either. We shall not try to sort out these intricacies fully here: the evidence we shall be talking about is either paradigmatically trace or paradigmatically predictive and the need to classify these more complicated or intermediate cases won't arise.

⁷ Sanchirico (2001), p. 1286.

any incentive to refrain from opening an umbrella, since those who scan positive will be fined come what may and those who scan negative will not be fined come what may:

Individuals deemed unlikely to open umbrellas on such days escape penalty regardless of their current behavior; they feel free to open their umbrellas for shade... High propensity individuals, on the other hand, find themselves in the classic position of being “damned if they do, damned if they don’t.”⁸

By contrast with its terrible incentive structure, this system may do quite well as far as the accuracy of its convictions go. People with the chemical have nothing to lose by opening the umbrella on sunny days since they will be getting a fine anyway. For all we have said it may well be that those receiving fine are invariably guilty of the act for which they are getting fined. But the incentive structure may well be problematic even given that accuracy rate.

This legal system is very extreme. Trace evidence is not used at all. By contrast, and as Sanchirico emphasizes, the calculation of incentive effects is quite different in cases when trace evidence and predictive evidence are *both* given a role. To illustrate, let us look at another quite simple case. Suppose the law is set up so that if the neighbor of x’s barn burns down without an eyewitness, and it can be shown that x had motive and a prior history of conviction for barn burning, that is sufficient for conviction.⁹ Let us further suppose that Jones knows all this. Jones is standing in front of his neighbor’s barn, canister and lighter in hand, and is deciding whether to set the barn ablaze. Suppose that Jones knows there will be no eyewitnesses and he has obvious motive – his neighbor got him fired just recently – and that he has a prior conviction for barn burning that is easily accessed. The case is crucially different to the weather scenario. Jones has a choice between burning and not burning. He knows that if he burns it is extremely likely he will be convicted of barn burning that night. He knows that if he doesn’t burn it is extremely unlikely that he will be convicted of barn burning that night. After all, if he doesn’t burn the barn that night it is pretty unlikely that anyone else will burn down the barn that night and that would be needed for being convicted of burning down the barn that night.

⁸ Sanchirico (2001), p. 1287, 1266. Notice that the generic association between the chemical and umbrella opening may be somewhat disrupted by the newly instituted legal system,. For example, there is (for some) a new pleasure – breaking the law with impunity – that may motivate people without the chemical to put up their umbrellas, which the doomed citizens may be too depressed to both putting up umbrellas on sunny days. Sanchiro’s set up is slightly different: In his presentation, “historical data on each individual indicating with reasonable accuracy whether she is prone to open umbrellas on sunny days” mean that “The individual knows that the legal consequences she faces on a sunny day are a forgone conclusion” (2001, p. 1287). But this raises tricky issues, since informing someone of predictions about their behavior may influence their behavior. The issues raised here are somewhat similar to those attending the deliverances of the predictor in the famous Newcomb problem. We have tried to bypass them by speaking of a generic association rather than an individualized one.

⁹ We can imagine it is a community which does not have the resources for collecting fingerprint evidence and DNA evidence etc.

Why is this case so different? The answer is clear enough. Even when there are no eyewitnesses, this is a system that uses trace evidence as part of the basis for conviction, and the presence or absence of that trace evidence is probabilistically highly dependent on whether Jones sets the barn ablaze. The fact that there are the charred remains of a barn demonstrating destruction of the barn by fire is, in this legal system, a crucial part of the case against barn burners. If the barn is not burnt down that night, Jones will not be convicted of burning it down that night. Notice that the trace evidence is not individualized: that there are charred remains of that barn the next morning is not a fact about Jones. But as we have seen, trace evidence need not be individualized.

When trace evidence and predictive evidence are both admissible, the calculation of the incentive effects will often be considerably more complicated. Compare system (a) which requires an eyewitness and a burnt barn for conviction with system (b) which requires a burnt barn and either an eyewitness or motive plus prior conviction. System (b) will provide extra disincentive to Jones. But system (b) will also provide incentive to those who wish to harm Jones. If Smith knows that he can burn down the barn with no eyewitnesses, gets no joy from barn burning but wants Jones in jail, then in system (a) he will not burn down x's barn but under system (b) will be highly motivated to do so. Of course it is rare that one knows that there are no eyewitnesses. To the extent that the presence of an eyewitness is typically uncertain, cases like the one where Smith burns to frame Jones may be predicted to be rare even under system (b). What bears emphasis then is that matters are far from simple when various combinations of trace and predictive evidence are both in play.¹⁰

3. The Blue Bus Case

Let us turn to the classic blue bus case, which has been used by legal scholars to raise a puzzle about the propriety of using statistical evidence as the basis for determining legal responsibility.

¹¹ Here is Enoch and Fisher's presentation of the puzzle (there is a very similar presentation in Enoch, Spectre and Fisher):

One starting point for the statistical evidence debate is the classic Blue Bus hypothetical, which is a variant of *Smith v. Rapid Transit, Inc.*, a seminal case in modern evidence law. The hypothetical consists of two cases. In both cases, a runaway bus injures the plaintiff, and the case goes to trial against the eponymous bus company. In the first case, the evidence includes eyewitness testimony that one of the Blue Bus Company's buses caused the injury. The witness, however, is imperfectly reliable. To

¹⁰ Sanchirico is generally very clear about this.

¹¹ The first application of the blue bus scenario to the statistical evidence issue was Judith Jarvis Thomson (1986).

illustrate, let us assume her to be 70% reliable in such circumstances.¹² In the second case, however, there is no eyewitness to the accident. Instead, the plaintiff seeks to introduce evidence about the Blue Bus Company's market share in the area where the accident took place. The uncontested market share data show that the Blue Bus Company owns 70% of the buses in the relevant area. This, the plaintiff argues, shows that it is more likely that one of the Blue Bus Company's vehicles was involved in the accident, because Blue Bus is the largest bus company in the area with the greatest number of buses on the road. Even though the evidence in both cases may be of equal probative value, our responses to the two cases are very different.

The puzzle, as standardly presented, is to justify the differential reaction between the two cases and, relatedly to explain what would be wrong about a civil court system in which, in the absence of eyewitnesses, it would be fair game to find the Blue Bus company liable on the grounds of market share. As noted earlier, the Enoch papers pursue two lines of thought here, one incentive based, one sensitivity based. We shall look at the sensitivity issue in the next section – it is the incentive based reasoning that concerns us here. Invoking Sanchirico, Enoch and Fisher say:

Ideally, in order to generate efficient incentives, we would want the actor to know that the likelihood of his being punished strongly depends on whether or not he decides to break the law here and now. The weaker this dependence, the weaker the incentive provided by the law to not engage in this specific criminal behavior. Thus, admitting character evidence at the trial stage would be counterproductive in terms of incentives. The prohibition on character evidence promotes deterrence by avoiding a decrease in the marginal cost of engaging in criminal behavior (2015, 582).

And later:

¹² One has to be careful here about what "70 per cent reliable" amounts to. Suppose it worked like this: Conditional on seeing a blue bus it is 70 per cent likely that the person will say "Blue bus" and 30 per cent likely that the person will say "Red bus" and conditional on seeing a red bus it is 70 per cent likely they will say "Red bus" and 30 per cent likely that they will say "Blue bus". To think we should then be 70 per cent confident that a red bus was responsible upon hearing "Red Bus" would be to commit the base rate fallacy. Perhaps Enoch et al are thinking of a case where market share is split or even more simply that it is .7 probable in the circumstances that the particular eyewitness report is correct. (If the former then better to be explicit about that. If the latter then it is best to avoid talk of reliability here since that way of presenting things encourages base rate fallacious thinking.)

Sanchirico's analysis can also be employed in the Blue Bus context: if statistical evidence regarding the 70% market share of the Blue Bus Company were admissible at trial, deterrence would be undermined. This is due to the fact that the Blue Bus Company's expected cost of engaging in negligent behavior is a function of the difference between the probability that liability will be imposed given negligence and the probability that liability will be imposed given engagement in the socially desirable behavior. Admitting the market share statistical evidence would enhance the probability of liability in the latter type of cases. In other words, introducing statistical evidence at trial (ex post) would lower the marginal cost of negligent behavior for the Blue Bus Company, thereby dampening its incentives to take the necessary precautions or to engage in the desirable level of activity (ex ante). At the same time, the Red Bus Company—holding only 30% of the market share—will also be disincentivized to adopt the socially optimal precautions or activity level so as to prevent the occurrence of negligent accidents, because introduction of the statistical evidence will lower the prospects that it will be held liable for such accidents." (2015, 583.).

Their reasoning about the blue bus case is incorrect. To make things vivid, let us imagine two contrasting legal systems. In system (a) if it is shown that a bus from town hit a car or person and neither company owns up to having done it, neither bus company pays up unless there is eyewitness evidence telling against one of the companies. In system (b) there are two ways that a bus company pays up when it is shown that a bus from town hit a car or person and neither company owns up to having done it: either there is eyewitness evidence telling against the company or the company has a 2/3 or greater market share of the buses on the road in that town. There are of course myriad other possible legal systems but it will be instructive for now to compare just these two.

Imagine a blue bus driver that is a shareholder in the blue bus company driving along the street one night. The bus driver sees that he can knock into a car and knows that as there are no passengers, there is no chance of eyewitnesses and no chances of getting hurt. He knows that if he hits a car there will be tell-tale signs of a collision by a bus rather than another vehicle but no distinctive features of the dent that would point to the blue bus company rather than the red bus company. (Such assumptions can be relaxed a bit, making various computations a little more complicated but without disrupting the main contrasts)

Imagine system (a) is in effect. Assuming zero probability of eyewitnesses, in system (a) the probability of the company being found liable for hitting that car conditional on hitting is zero and conditional on not hitting is zero. So if the driver derives some utility from hitting the car he will hit and if the driver derives no utility from hitting but derives disutility from taking a certain amount of extra care to guarantee – as opposed to merely making it probable – that the driver does not hit, then the driver will not be incentivized to take extra care. But in system (b) the probability of the blue bus company being found liable for hitting that car conditional on hitting is

one and conditional on not hitting is whatever small probability there is of another bus hitting that car that night. In other words, in system (b) the probability of a blue bus driver's company being found liable depends in this case much more strongly on whether one commits the dubious act than in the no statistical evidence regimen. Meanwhile, in the case of a red bus driver in the same circumstance the two regimes make no difference to the likelihood of a red bus driver's being found liable conditional on hitting the car. Where the probability of an eyewitness is zero, the probability of liability is zero either way.

Imagine instead there is a 20 percent chance of eyewitnesses – it is a bit foggy and there are no eyewitnesses in sight but the lack of eyewitnesses is by no means guaranteed. (To keep things simple imagine eyewitnesses are perfectly accurate – though the main points will not depend on this simplifying assumption, as any reader can see by redoing calculations on assumptions of less than perfect accuracy.). Then in system (a) the probability of being found liable of hitting that car on that night conditional on hitting is .2 (ignoring extra tiny probabilities having to do with a second blue bus hitting the car for a second time with an eyewitness) and the probability of being found liable conditional on not hitting is whatever tiny conditional probability there is of another blue bus hitting that car that night with an eyewitness. Meanwhile in system (b) the probability of the blue bus company being found liable of hitting that car conditional on hitting conditional on hitting is 1 (since whether or not one there is an eyewitness it will be found liable, at least ignoring tiny probabilities of the car being hit a second time by a red bus with an eyewitness etc) and the probability of the blue bus company being found liable of hitting that car conditional on not hitting is whatever tiny conditional probability there is of a red bus hitting that car that night without a eyewitness plus whatever very small conditional probability there is of a different blue bus hitting that car that night. Meanwhile, as before, the two regimes make no difference to the likelihoods of liability for a red bus driver.

Of course things are different if an individual bus driver does not care about their company and only worry about hitting a car insofar as an eyewitness identifies *them* as a driver. It is helpful then to see that the two regimes can have somewhat similar effects when it comes to large scale planning by the company directors (though the nature of the effects depends crucially on the details). Suppose the bus companies can invest in very expensive tires. The blue bus company expects to hit about 20 cars in the year in situations without eyewitness and about 10 with the super expensive ones. In regime (a) the blue bus company will expect to pay out 20 times with the current tires and about 10 times without, a difference of 5 payouts. Conditional on regime (b) the blue bus company will expect to pay out 24 times with the current tires and 14 times without, a difference of 10. So regime (b) increases the incentive for buying the expensive tires.

Regime (b)—in which statistical evidence is used—adds extra incentive for the blue bus company to avoid hitting buses in all sorts of scenarios, both at the level of an individual bus driver's decision and at the level of company planning. But this is not what Enoch et al say. They think that statistical evidence regime weakens the incentive not to do the bad thing, and that just isn't right.

We have a speculation as to the source of the confusion. A kind of case that figures prominently in the two papers is a gatecrasher case.¹³ To fix ideas, consider the following vignette:

George gets to the barrier of the club and is deciding whether to gatecrash. There is a big crowd behind him and so he has no choice but to enter the club. His options are to gatecrash or not to gatecrash. He knows that once a week the police raid the club looking for gatecrashers. Imagine two regimes: Regime (a): The club issues a ticket to those who pay to get in, and when the police raid they give fines to those without a ticket. Regime (b) The club collects cash at the door, does not issue a ticket and the police decide whether to fine based on the difference between the number of people indicated by the cash collected and the number of people inside. If the number inside is more than quadruple the number who have paid to get in, the police give everyone inside a fine. Since gatecrashing is so common, they end up giving out a fine half the time and George knows this.

The gatecrasher case is radically different from the blue bus case. Enoch et al are right to think that when the choice of regimes is like that in the Gatecrasher case, the incentive to buy a ticket is strengthened when trace evidence rather than statistical evidence is used as a basis for conviction. In regime (a) the chances of getting fined are more or less zero conditional on buying a ticket (we say 'more or less' to control for lost tickets) and one in seven conditional on not buying a ticket. In regime (b) the chances of getting fined are one in fourteen whether or not one hands over the cash, and so the incentive to buy a ticket is dampened. But their extrapolation from these dynamics to something analogous in the blue bus case is incorrect.

There is a big difference between the gatecrasher case and the blue bus case. Why is this? It is instructive to look at the cases through the lens of putative trace evidence. In the gatecrasher case putative trace evidence is used as a basis for fining George during a raid. What is it? The putative trace evidence is George's presence inside the club at the time of the raid! But notice that this putative trace evidence will be there whether or not George pays for a ticket.¹⁴ So there is no incentive (apart from internalized shame etc.) for him to buy a ticket. (In the slightly different scenario where George has the option of not going to the club, regime (b), while making it pointless for club-goers to pay, will disincentive club going. Not going to the club will ensure that the putative trace evidence – one's presence at the club – will not be found during a raid.) In short, the problematic act of entering without paying makes no difference (vis a vis

¹³ The gatecrasher case was introduced into the literature by Cohen (1977)

¹⁴ Consider similarly in well known prison riot case from Nesson (1979). A prisoner participates in a one hundred person riot where ninety-nine are known to be involved in killing anon guard, though it is not known which prisoner did not participate. Here the trace evidence of the guard's murdered body would be there whether or not the non-participant had participated.

entering and paying) to the chance of being fined in regime (b) but makes a big difference to the chance of being fined in regime (a).

Consider next the Blue Bus scenario. What is the putative trace evidence when the blue bus hits a car on a certain night? The putative trace evidence is the bus-collision-indicative dent on the car. But crucially, the expectation of that putative trace evidence conditional on the problematic act is vastly different to its expectation conditional on not performing the problematic act. If you are a blue bus driver then conditional on not hitting a certain car, the probability of a bus-indicative dent being on that car the next morning is very small (it's a matter of the probability of another bus hitting that car), but the probability of a bus-indicative dent being on that car conditional on hitting that car is high (on our simplifying assumption that collision by a bus guarantees a bus indicative dent the probability is 1). In the blue bus case, but not in the gatecrasher case, the likelihood of the putative trace evidence is heavily dependent on whether or not the agent performs the problematic act; this makes a difference. That the dent is putative trace evidence of a blue bus hitting the car is completely uncontroversial. The likelihood of a blue bus having hit that car conditional on the presence of a bus-collision indicative dent is obviously higher than conditional on the absence of such a dent.¹⁵

What the statistical evidence regime does is magnify the upshot of the putative trace evidence (which in this case is bona fide trace evidence) when a bus collides. With or without the statistical evidence regime, the blue company will pay up when there is the putative trace evidence of a bus collision in combination with the putative trace evidence of a blue bus collision reporting eyewitness. But in the statistical evidence regime, a blue bus hitting a car guarantees putative trace evidence sufficient for rendering the blue bus company liable (since with or without eyewitnesses there will be evidence sufficient for liability). Meanwhile if the blue bus does not hit the car then in neither regime is it likely at all that the blue bus company will be found liable for hitting that car that night since conditional on that blue bus not hitting that car that night it is not at all likely that another bus will. Thus there is a significant difference in putative trace evidence structure between the gatecrasher case and the blue bus case. We suspect that the reason that Enoch et al do not give the cases separate treatment is that they do not spot the difference. Notice moreover that there is nothing in the blue bus case that makes trouble for Sanchirico's main points, since in the statistical evidence regime the blue bus case is one where trace evidence and the predictive evidence provided by market share work in combination.

In saying all this we do not wish to claim that something like the statistical evidence regime driven by market share is a good idea for any legal system devised to handle bus collisions. Even granting that such a regime might well provide extra incentive to blue bus company

¹⁵ The Uviller/Sanchirico distinction between trace evidence and predictive evidence features in Dahlmann's discussion of incentives and statistical evidence. But having made the abstract distinction between "cases with only naked statistical evidence and cases with a mix of naked statistical evidence and trace evidence", he classifies the blue bus case as one where "there is only naked statistical evidence". He also remarks that "naked statistical evidence is predictive". His discussion thus also obscures the fundamental contrast between the gatecrasher case and the blue bus case.

directors to buy the better tires and provide extra incentive for a blue bus driver to avoid a collision in the scenario described there is plenty of reason to be dubious about the regime.

In the particular setting in question, accuracy concerns are legitimate. Let us say that a finding of liability against a bus company for hitting a car is accurate when and only that bus company is actually causally responsible for hitting the car.¹⁶ When eyewitness evidence is required, then if the probability of any eyewitness report being true is far greater than two thirds, there will be far less inaccuracy when an eyewitness is required than when statistical evidence is a sufficient basis. But as Enoch et al are well aware, if we move to a scenario where the market share of the blue bus company is much higher – say 98 per cent – and where the occurrence of lying or perceptual error among eyewitnesses is at least two percent, then there will be no obvious gain in accuracy when finding liability against the red bus company on the basis of an eyewitness report then when finding liability against the blue bus company in the absence of eyewitnesses. (That is the sort of consideration that prompted the Blue Bus puzzle in the first place.) Moreover, unaccountability for dodgy deeds is also undesirable. A bus company is unaccountable for hitting the car when it does so but there is no judgment against it. A statistical evidence regime can substantially reduce overall unaccountability for dodgy deeds. In short it is not a straightforward matter to show, in a general way, that the statistical evidence regime is bad in the Blue Bus-style scenario simply by focusing on the badness of inaccuracy coupled with the badness of unaccountable car destroying behavior.

Nevertheless there are plenty of reasons to be concerned about the statistical evidence regime in the blue bus case, many of which are incentive based.

One obvious concern about the simple statistical evidence regime is that it provides a great deal more incentive for the blue bus company to avoid collisions than the red bus company. Imagine that the red bus company directors are deciding whether to buy expensive tires. Suppose they expect 6 collisions with the current tires and 3 without. Moving from regime (a) to (b), while providing additional incentive for the blue bus company to upgrade, makes no difference to the expected payouts for the red bus company. This raises concerns of fairness. Isn't it unfair that to have a system in which, when the blue bus company is hitting 70 per cent of the cars, it has to shoulder the burden of paying a much higher proportion than that to plaintiffs? System (a) removes that unfairness though at the cost of making the overall liability for bus companies much lower, thus placing the economic burden on the insurance companies of car owners (and indirectly on car owners themselves if it results in increased premiums). But there are other

¹⁶ It is important to contrast this with more global notions of accuracy. Let us say collision liability findings against the blue bus company are globally accurate over a period to the extent that the gap between the number of such findings and the number of actual collisions during that period is small. The number of liabilities during the year may match the number of cars hit by the company even if there are local failures of accuracy. And finding of liability may be perfectly locally accurate but the number of findings may fall far short of the total collisions with a consequently poor global accuracy score. The relative weights of local and global accuracy in civil settings is a good question.

systems – notably where the liability is split 70-30 in the absence of eyewitnesses – which remove the unfairness while retaining the overall burden on bus companies. (Of course this kind of system will not be practical in a situation where collisions are either by the Blue Bus company or one of many long distance coach companies coming through town -- where the proportion of blue buses on the roads in town compared to the combined total other myriad companies is around 70-30.¹⁷)

This unfairness also has economic correlates. With this last mentioned system in place, the red bus company has more incentive viz a viz the other regimes to buy the better tires. Moreover, if the rule is that the company with the market share over two thirds is liable when there are no eyewitness, then (as Enoch and al. point out in a footnote), companies with a slightly less than two thirds market share will be disincentivized from expanding. Meanwhile companies with a greater market share will be incentivized either to contract or to fragment into a few legal entities – say a Dark Blue bus company and a Light Blue bus company – in order to bypass liability. (Shared liability across the bus companies in cases of liability would mean that no percentage serves as a special threshold in the way two thirds does in system just alluded to, and would mean there is no special advantage of the sort just described that is achieved by fragmentation)

There is also the risk of diminished quality of evidence gathering – if market share is a sufficient basis for a finding of liability, then there is little incentive for the insurance company of the plaintiff to track down eyewitnesses since they care about the payout from the bus industry far more than which company does the paying. Notice that this downside persists even when we have a regime with shared liability across bus companies and so would have to be in part controlled for by requiring a certain amount of evidence gathering.

Note also that there is extra incentive for the red bus company to be malicious. If the blue bus company is the competition, it may serve their interests to harm it financially. In a case where there is a known absence of eyewitnesses, such harm can be accomplished by hitting a car. We do not wish to make too much of this last point however – obviously a known absence of eyewitnesses is extremely rare (not to mention the fear of the bus driver's getting hurt) , and unless the utility generated by harming the blue bus company is enormous, a decent risk of eyewitness or injury will typically suffice to disincentivize deliberate collisions.¹⁸

¹⁷ Note that the original Smith vs Rapid Transit case had more of these elements.

¹⁸ Dahlmann remarks that the statistical evidence regime “gives companies with a minor market share an incentive to be opportunistic and act negligently, since the consequences of such behavior is carried by the company with the major market share.” (2020, 176.) He then suggests that “it is reasonable to assume that this negative effect “is stronger than the positive effect of the incentive structure of the company with the major market share, and this means that the net-effect is negative.” (2020, 176) In support of this verdict of comparative strength he suggests that the incentive for the blue bus company might not be very strong since the greater market share company would “have to pay for all the accidents caused by the negligence” of the lesser market share company (2020, 176). This is far too quick. As we noted above, the minor company's incentive to opportunistically try to damage the majority company will be significantly dampened by the fact that the absence of eyewitnesses is not known in advance. So the incentive for opportunistic interference may be small. Meanwhile, depending on the frequency of eyewitnesses, the incentive for the blue bus company to avoid accidents may be *significantly* greater in a statistical evidence

It is also worth remembering how artificial regime (b) is. Market share is not the only kind of predictive evidence that has a probabilistic bearing on which bus company was responsible. Suppose blue bus drivers are far better trained than red bus drivers. Why should this be rendered irrelevant? System (b) diminishes the incentive to demonstrably offer high quality training compared with a system that allowed the quality bus driver training to count as relevant. And what if (perhaps because of differential training, perhaps because of better tires or perhaps for reasons unknown) the ratio of eyewitnessed blue bus collisions to red bus collisions is far closer to 50-50 than one might expect from market share? Shouldn't this be allowed to be relevant if market share evidence is? Suppose to accommodate this we allowed any potentially relevant predictive evidence into the courtroom. Then we would be placing huge demands on a jury to process the joint probabilistic impact of the predictive evidence, something which human beings are not good at (Enoch et al. gesture at a range of practical concerns like this themselves). Of course one could have a specially trained expert brought in that did the job of making some judgment based on evidence available, but the time, effort, and expense of bringing in experts in a case where a bus hit a car might make no economic sense.

In short, it is pretty easy to see that regime (b) is a pretty bad idea. But Enoch and al's diagnosis of why it is a bad idea is incorrect. And as we shall see, it is by no means clear that systematic deployment of statistical evidence as a basis for finding liability is a bad idea in all areas of the law.

2. Conditionals and Statistical Evidence.

regime than in one where eyewitness evidence is required. Suppose that the blue bus company expected to hit 100 cars without a tire upgrade, 50 with, and expected about half the collisions to have eyewitnesses. The result of shifting from regime (a) to (b) will be one of placing additional burden on the blue bus company. But no matter how many cars the minority company is expected to hit, this makes no difference to the differential in expected outlay for regime (a) as between new and old tires and no difference to the differential of expected outlay in regime (b), and thus no difference to the extra incentive provided by (b) to get tire upgrades, at least on one natural measure. If the red bus company is expected to hit two cars, then upgrading changes the expected payout from 50 to 25 in regime (a), a differential of 25, and 102 to 52 in regime (b) a differential of 50. So if the tires are only good for a year and cost the equivalent of x payouts, they are a good idea in regime (b) when x is less than 50, but are only good in regime (a) when x is less than 25. Suppose the red bus company is instead expected to hit fifty cars. Then upgrading changes the expected payout from 50 to 25 in regime (a), a differential of 25 (b) and from 125 to 75 in regime (b) a differential of 50. The facts about whether a tire upgrade equivalent to x payouts is a good idea have not changed. Greater collisions by the minority company increase the burden on the blue company, but do not change whether a tire upgrade equivalent to x payouts is a good idea. Thus Dahlmann's suggestion that incentives for the blue to avoid accidents are significantly dampened by the need to cover the costs of the red company's accidents is at best highly misleading. More generally, we simply can't make the kind of sweeping comparisons about bus company incentives prefaced by Dahlmann's "it is reasonable to assume" since so much will depend on details about context.

2.1 Sensitivity and Blue Bus Intuitions.

Enoch et al. introduce sensitivity in the following way:

Sensitivity: S's belief that p is sensitive =df. Had it not been the case that p, S would (most probably) not have believed that p (212, 204).

Counterfactuals associated with sensitivity play two key roles in Enoch et al.'s discussion

First, they use such counterfactuals to explain the "overwhelmingly strong and common intuition among practitioners and scholars alike that there is something suspicious" (2012, 198) about finding the blue bus company liable on the basis of market share consideration. However, this does not, they think, provide a normative story as to why the legal system should refrain from finding liability on such grounds. By way of normative justification, they appeal to an incentives story, but argue that the "very same counterfactuals the epistemological account relies on" (2012, 212) will play a starring role in explaining why a legal system that routinely judges companies liable on the basis of statistical evidence will have bad incentivizing effects.

In arguing for a sensitivity based diagnosis of the relevant intuitions, they have the following to say about the blue bus case (where 'the second scenario' refers to case where liability is based on market share rather than an eyewitness report)::

Things are different, though, if we base our finding solely on statistical evidence, as we do in the second scenario. In that scenario, we find against the Blue Bus Company solely on the basis of its market share. Now, had it not been one of its buses that caused the harm, nothing would have been different regarding the market shares. In such a hypothetical scenario, the Blue Bus Company still owns 70 percent of the buses; it is just that the bus that causes the harm is no longer one of its buses (rather, it is a Red Bus bus). In such a case, we would still have the exact same statistical evidence available to us. So in that case, too, we would have found the Blue Bus Company liable. So by relying on statistical evidence, we render our findings insensitive (212, 206).

Here, they are eliding over an important distinction that, as we shall see, they are more explicit about elsewhere. Let us distinguish two kinds of conditionals concerning a blue bus collision. One deploys an antecedent that includes the fact of a bus collision, one that does not:

Type one

1a. If a blue bus hadn't hit that car that night the judge would (most probably) not have believed that the blue bus company most probably hit that car that night

1b. If that blue bus hadn't had that car that night the judge would (most probably) not have believed that the blue bus company most probably hit that car that night

1c. If a blue bus company bus hadn't hit that car that night then the judge (most probably) would not have found the blue bus company liable.

Type two

2a. If a blue bus hadn't hit that car that night and a bus did hit that car that night, the judge (most probably) wouldn't have believed that the blue bus company most hit that car that night

2b. If that blue bus hadn't hit that car that night and a bus did hit that car that night, the judge (most probably) wouldn't have believed that the blue bus company most hit that car that night

2c. If a blue bus company bus hadn't hit that car that night and a bus did hit that car that night, then the judge (most probably) wouldn't have found the blue bus company liable.

Let us type one "plain conditionals" and call type two "bad act fixed" conditionals. What Enoch et al. are seeing is that when the legal regime is something like (b), bad act fixed conditionals like those listed above are false, including ones that have findings of liability in their consequent. (Of course even in regime (a) the bad act fixed conditionals concerning the judge's belief may be false but that will be less concerning from a legal point of view as such beliefs wouldn't have legal consequences). And in Enoch and Spectre's recent (2019) they are particularly explicit that it is the bad act fixed conditionals better captures what they care about:

So if a Sensitivity-style story is to have any chance at all, the fact that the crime (or tort etc.) has occurred must be held fixed, as part of the counterfactual's antecedent (2019, 191).

It is, however, worth taking some care with the question of which beliefs are supposed to be insensitive here. Suppose that a blue bus hits the car, there are no eyewitnesses, but on account of market share, the judge believes a blue bus company bus caused that car to get hit that night. That passes sensitivity: if a blue bus company bus hadn't hit that car, then most probably no bus would and in that case the judge would not have believed that a blue bus hit the car that night. Similarly for a finding of liability – if a blue bus had not hit that car that night, most probably no bus would have and no finding of liability for hitting that car that night would have been made against the blue bus company.

Now of course since the judge believes a blue bus hit that car that night, the judge (assuming logical competence) also believes that either no bus hit that car that night or a blue bus did. Here is the sensitivity test for this belief:

If it hadn't have been that (either no bus hit that car that night or a blue bus company did) then the judge would most probably have not believed that (either no bus hit that car that night or a blue bus company did)

The sensitivity test is failed in a case where there are no eyewitnesses around. For the disjunction to be false it has to be that a bus that isn't a blue bus hits the car. But if that had happened the judge would most probably have believed on market share grounds that a blue bus hit the car that night. So in a case where a blue bus hits a car with no eyewitnesses around, it is this disjunctive content that fails the sensitivity test. Thus, if we think insensitivity is driving our concerns in the blue bus scenario (where a blue bus in fact caused the accident), it looks like it is because we recognize that the judge and other relevant legal protagonists fail the sensitivity test for the disjunctive belief that either no bus hit the car or a blue bus hit the car. But we submit that it would be quite odd for intuitive concerns about a regime that allowed findings of liability based on market share to be driven by sensitivity failures for somewhat baroque disjunctive propositions.

That said, even if the intuitions about the blue bus are not being driven by the concern that the judge has an insensitive belief, it is still open to Enoch et al. to argue that such intuitions are being driven directly by recognition of the falsity of bad act fixed conditionals. But even here we are suspicious. Juxtapose the blue bus case with the following case:

Ford Bumpers

The Blue Bus and the Red Bus Companies have the same number of buses on the roads in town. The entire fleet of the Blue Bus has Ford bumpers. Only 5 percent of the Red Bus Fleet has Ford bumpers, the rest having Chrysler bumpers. A bus collides with a car with no eyewitnesses. When a bus hits a car, it typically leaves an imprint showing

the brand of the bumper. A bus has collided with a car. There are no eyewitness but there is an imprint of a Ford bumper on the car that was hit.

Suppose it is suggested in this case that it is fine to find the Blue Bus company liable on the grounds that the ratio of Ford Bumpers between the two companies is 20-1. Notice in this setting that the conditional

If a blue bus had not hit that car that night and a bus had, the judge (most probably) would not have believed that the Blue Bus company did it

is *true*. After all, if a blue bus had not hit and a bus had, the bus would probably have left a Chrysler imprint. Thus if the bad act fixed conditional is really the diagnostic for the intuition that finding the blue bus company liable in the original case is bad, then we should predict that scholars and practitioners would not balk at finding the blue bus company liable in this variant case. We find it somewhat hard to believe that the pattern of intuitions will break this way, though we welcome an investigation of the matter.

It is instructive to see why a certain retreat will not do here. According to the retreat we have in mind, the bad act fixed conditional is not the real test for whether we intuit a finding of liability unacceptable, but instead the right test is given by a conditional that holds all the trace evidence fixed:

Fixed Trace Evidence Conditionals:

If the blue bus company had not done it and all trace evidence was the same, then the judge would (most probably) not have believed that the blue bus company did it.

Sure enough, this is false. Moreover, this kind of conditional is at least hinted at in Enoch and Spectre recent response to Pardo where they write:

What seems needed is to hold fixed the occurrence of the trial, and ask whether you would have acquitted had there been a trial, and had *everything else been pretty much as it actually is* (our emphasis), except that the defendant had been innocent. (2019, 191)

Sure enough, the proposition that the Blue Bus company did it fails *this* test in the case of Ford Bumpers. But pretty much any case of judged liability will fail this test. Suppose for example we find liability on the basis of an eyewitness report that a blue bus company bus hit the car. If the blue bus company had not hit and all the trace evidence had been the same (which requires holding fixed the fact of the report), the judge would still have believed the blue bus company did it. So the Trace Evidence Fixed Conditionals do not explain differential reactions to the case

where liability is judged on the basis of market share and cases where it is judged on the basis of eyewitness reports.

In sum, we do not think that Enoch et al have provided a convincing story about the intuitive reactions of scholars to the blue bus case. (We postpone further treatment of that issue to another time, since incentive structures are our main focus here).¹⁹ But what about their story about the incentivizing role of conditionals? It is to this issue that we now turn.

2.2 Conditionals and Incentives

Enoch et. al think that sensitivity considerations explain our intuitions about the blue bus case but that sensitivity failures cannot by themselves justify a legal system in which statistical features like market share are used to judge guilt or liability. (As they think about it, this would involve a kind of “epistemological fetishism” (2012, 213).) Nevertheless they think that the conditionals that figure in their sensitivity story also play a starring role in explaining the relevant incentive structures, and that these incentive structures are a legitimate basis for legal abhorrence to convictions, liability finding, etc. based on statistical evidence. Let us see how they attempt to connect the sensitivity story to the instrumental incentive based considerations:

Think about incentives again, say, in the case of John, who is deliberating about whether or not to purchase a ticket. He is now thinking in terms of conditionals, things like “if I crash the gates, they will punish me. If I do not, they will not.” And typically, when at a point in time some such conditionals are true, at a later point in time (some of) the very same facts are captured by counterfactuals, or subjunctive conditionals. Suppose that John proceeds to crash the gates. Then his conditional “if I do not crash the gates, they will not punish me” captures the fact that we can now—say, when John is on trial—capture with “had he not crashed the gates, we would not have punished him.” And this counterfactual should sound familiar to you: it is the relevant instance of Sensitivity. In other words, though the epistemological story is not itself of legal value, and though the instrumental story that is of legal value is not itself epistemologically respectable, both of them still stem from the same source, Sensitivity-style counterfactuals.(2012, 220)

And then a little later they write that their incentive story

. . . relies on the truth of relevant counterfactuals, indeed the same counterfactuals the epistemological story relies on (2012, 221).

¹⁹ Nor do we want to get into the relation between sensitivity and knowledge -- Enoch et al. think that sensitivity drives our intuitions because we use sensitivity to test for knowledge. We disagree about the relation of sensitivity to knowledge but a discussion of that issue would take us too far afield.

A few points bear emphasis.

As we have already noted, the gatecrasher case and the blue bus case are sharply different. One aspect of this difference is that there are marked contrasts in judgments by the protagonists of those cases concerning conditionals. In a gatecrasher case where one knows that a raid is coming and that one will be fined come what may, one will think: 'If I pay I will get a fine. If I don't pay, I will get a fine'.²⁰ By contrast, suppose one knows there are no eyewitnesses and that a market share conviction regime is in place. As a blue bus driver one will *not* think "If I hit that car my company will get found liable for hitting that car tonight and if I don't hit that car my company my company will get found liable for hitting that car tonight". After all, one will think that if one doesn't hit the car then most likely no bus will. If there is anything wrong with the market share conviction regime, it is not something that reveals itself in "damned if we do, damned if we don't", thoughts by drivers or market share directors.

Granted, Enoch et. al concede themselves that "in some cases incentive structures work in a somewhat more complicated way" (2012, 218), mentioning that in the blue bus scenario, the key deliberations are more likely to be done by organizations than by specific agents. But they don't play up the most basic contrast between the gatecrasher scenario and the blue bus scenario, namely that in the latter, the driver will not think "If I don't hit the company will still get found liable" and the organization won't think "If we hit a lot less cars this year we will be found liable for just as many collisions."

Interestingly, they say a little later that

sometimes (in cases unlike the gatecrasher case) the deliberating would-be perpetrator knows that if he chooses not to commit the crime, he knows no crime will be committed at all," (2012, 218-9)

They recognize that the incentive structure will be different here, though do not notice a connection between the theme of these remarks and the blue bus case.²¹ In this connection, it is

²⁰ Of course even in the gatecrasher case things are different if one doesn't know whether raid is coming. In that case one will not affirm 'If I don't pay I will get fined'. What really matter is that the probability of getting fined conditional on paying is the same as conditional on not paying. We say more about the relative importance of conditional probabilities vis a vis conditionals in the paragraphs that follow.. We also note in passing that these are indicative conditionals, not counterfactual conditionals. Enoch et al. are assuming that the proposition expressed by such indicative conditionals as 'If I don't pay I will get a fine' are truth conditionally equivalent to the propositions that are later expressed by counterfactual conditionals such as 'If I hadn't paid, I would have gotten a fine.' This assumption is tendentious, but we shall not bear down on it here.

²¹ The theme of the quoted passage reemerges in Enoch and Fisher. say that there may be justification for using statistical evidence as a basis for conviction when it comes to acts that occur within the personal relationships, since "acts in the context of a personal relationship, the deliberating

important to realize that even when the would-be perpetrator does not *know* whether the relevant bad act will be committed if they do not do it, differential conditional probabilities can similarly make for a very different incentive structure to the gatecrasher case. And in the blue bus case there are differential probabilities: the likelihood of the blue bus company's being found liable for hitting a car conditional on the blue bus company hitting is very different to the likelihood of the blue bus company's being found liable conditional on it not hitting.

In sum, Enoch et al.'s account is tailored to the fringe gatecrasher case where the 'damned if I do, damned if I don't' thought is correct. Such cases are exceptional and do not shed much light on various more routine cases -- like the blue bus case -- where the merits of judging liability based on statistical evidence are being debated.

Notice, moreover, that the bus driver, insofar as she engages in conditional reasoning, will evaluate plain conditionals not bad act fixed conditionals. The bus driver might find it natural to wonder what will happen if they knock into a certain car as opposed to what will happen if they don't. But the driver will not find it natural to wonder what will happen if they hit the car as opposed to what will happen if they don't but another, differently colored, bus does. Similarly, companies will be much more interested in 'How will things be if we run into half as many vehicles this year?' than in 'How will things be if we run into half as many vehicles this year as last year and the red bus company runs into enough extra vehicles to make up the difference?' As we have seen, it is bad act fixed conditionals that are supposed to generate our epistemological intuitions in the blue bus case. But these are not important to incentive-revealing reasoning at the level of the bus driver or the company.

It bears emphasis that what matters to the protagonists is the probability of certain sorts of putative trace evidence being produced. If a red bus driver is wondering about the consequences of hitting a car, what matters is not whether eyewitnesses are in fact around but the likelihood from that driver's point of view of eyewitnesses. The same point holds for any other sort of putative trace evidence that might have a differential impact on judgments of liability—what matters is the comparison between the bus driver's credences about such evidence emerging conditional on him hitting a car and his credences about such evidence emerging conditional on him not hitting a car. These credences will (in light of the relevant legal regime) inform the comparison between the bus driver's credence that the blue bus company will be found liable for damages conditional on him hitting a car and his credence that it will be

perpetrator typically knows that if he chooses not to commit the particular crime or take the injury-causing course of action, no such act will be inflicted on the potential victim. (2015, 608). They then quickly contrast crimes within a personal relationship with "the gatecrasher case, where the would-be perpetrator knows that even if she doesn't gatecrash the stadium, others will still do so. (2015, 608) They are wrong to think that this points to an important distinction between acts within the context of a personal relationship and acts that are not. After all, the market share regimen in the blue bus case is hardly a personal relationship setting, but that case that contrasts sharply and in relevantly similar ways with the gatecrasher case.

found liable conditional on him not hitting a car. These conditional credences are what matter for decision-theoretic purposes.

The above considerations indicate that if one wishes to think about the knock on effects for incentives of this or that legal regime, it is best to think about its effects not on the assertability of conditionals but instead on conditional probabilities. In thinking about incentives and disincentives to perform various acts as between legal regimes X and Y, we should not typically be looking at the outright assertability of such conditionals as “If I perform x, bad consequence y will eventuate”. Suppose in legal regime X it was 2 per cent likely (in a certain context) that one gets a month in prison conditional on drinking and driving and in legal regime Y it was .01 per cent likely”. Under neither regime can one assert “If I drink and drive I will get a month in prison”. But the incentive structures are quite different and the contrasting conditional probabilities of various outcomes conditional on various actions seems exactly the right way to probe the comparison of incentives.

When evaluating the knock on incentivizing effects of a legal system, the resultant conditional probabilities for protagonists reasoning about whether to engage in problematic behavior is obviously relevant. But also potentially relevant are the effects on conditional probabilities for protagonists reasoning about whether to engage in completely unproblematic behavior. Consider for example a legal regime in which, in an effort to reduce drugs in cars, there were routine searches of people driving to town. Suppose 20 per cent were searched. This may disincentivize innocent people from driving into town on the grounds they find searches unpleasant and humiliating and who reckon the probability of being searched conditional on driving into town to be high enough to discourage them from driving. Consider similarly the incentivizing effects viz a viz tire upgrades of a statistical law regime. The buying of better tires is unproblematic behavior, but the presence or absence of the statistical law regime in the blue bus case makes a difference to the expected benefits of doing it.²²

Even granting that ordinary conditionals are too blunt an instrument in this area, someone might suggest that, instead of focussing on conditional probabilities one might instead try to get the probabilities of conditionals to do the relevant explanatory work. This approach is not obviously bad, but will get one unnecessarily enmeshed in tricky philosophical issues. Suppose the police toss a coin at 11:00 PM to decide whether to raid. It's clear enough that conditional on attending the club, it's roughly .5 likely that one will witness a raid at that club. But what of the counterfactual ‘If one were to attend, one would witness a raid’? Suppose one in fact stays home. Inspired by David Lewis's (1973) approach to counterfactuals, Alan Hajek (ms and forthcoming) and others have held that counterfactuals with false antecedents whose consequents turn on a later coin flip are definitely false: The closest worlds contain a mix of heads and tails worlds and so it is false that at all the closest worlds where one attends the club

²² Considerations like this make us a little wary of Dahlmann's advice that we analyse incentive structure in terms of incentives to behave lawfully and incentives to behave unlawfully. Insofar as two regimes make a difference to whether innocent people drive to town, that is neither a difference in incentive to behave lawfully nor a difference in incentive to behave unlawfully.

it is raised and false that at all the closest worlds where one attends it is not raised. Meanwhile, there is a vast literature concerning limits on the extent to which conditional probabilities of indicative conditionals like “If I do such and such then such and such will happen” can line up with the conditional probabilities (and there are well known proofs that the correspondence cannot be perfect, see Lewis 1976, Stalnaker 1976 and Hajek 1989). Meanwhile others like Timothy Williamson (2020) have defended a material conditional treatment of the truth conditions of indicative conditionals that induces blatant mismatch between the conditional probability of a consequent on the antecedent and the likelihood of an indicative conditional’s being true. Meanwhile those who think with Dorr and Hawthorne (ms) that indicative conditionals have domains restricted to the epistemically possible will think indicatives of the form “If I do P, then Q’ will go vacuously true when it is epistemically impossible that one does P, but the corresponding counterfactual will not be vacuously true. (This explains why ‘I’m a monkey’s uncle’ is a routine enough consequent for indicative conditionals but “I would be a monkey’s uncle” is an eyebrow-raising consequent for counterfactual conditionals.) In short, if you want to go the route of assessing the probabilities of indicative and counterfactual conditionals in your story about the incentivizing effects of legal regimes, be our guest. But you will make life a lot easier for yourself if you stick to conditional probabilities.

Section Three: Toxic Torts

The hypothetical setting of homicides by epistemologists is an unconvincing case for giving statistical evidence a central role in criminal law or torts. A much more promising setting is that of toxic torts. Here is Black and Liliensfeldt’s characterization of such cases:

Toxic tort cases [are] those in which the plaintiff seeks compensation for harm allegedly caused by exposure to a substance that increases the risk of contracting a serious disease, but does not cause an immediately apparent response. These cases generally involve a period of latency or incubation prior to the onset of the disease. In most cases the increased risk of the disease does not diminish or dissipate, even with the cessation of exposure. (1985, 732)²³

Simply put, in a toxic tort case, in addition to establishing the defendant’s negligence, a plaintiff must establish both that a particular substance is toxic and that the toxicity of that substance

²³ For additional helpful discussion, see Harris (1986), p. 912.

(more likely than not) caused her injuries.²⁴ To establish the toxicity of the substance in question, plaintiffs often rely on epidemiological studies -- evidence that is purely predictive. To establish that the toxicity of the substance caused her injuries, the plaintiff most often provides expert testimony (called differential diagnoses or differential etiology) that other causes of her injuries are unlikely.²⁵

While a complete picture of how liability can be established in toxic tort cases is beyond the scope of this paper, it may be helpful to examine a leading case to see how things can look in practice. In *Merrell Dow Pharmaceuticals Inc v Havner*, the Texas Supreme Court ruled that the evidence presented by the plaintiff was insufficient as a matter of law to establish causation. In doing so, the court made clear what it thought would be sufficient for such a finding. First, in order to establish the toxicity of the drug in question, the plaintiff would have to establish that the drug increased her relative risk of injury by a factor of at least two. Specifically, the court wrote:

Assume that a condition naturally occurs in six out of 1,000 people even when they are not exposed to a certain drug. If studies of people who did take the drug show that nine out of 1,000 contracted the disease, it is still more likely than not that causes other than the drug were responsible for any given occurrence of the disease since it occurs in six out of 1,000 individuals anyway... However, if more than twelve out of 1,000 who take the drug contract the disease, then it may be statistically more likely than not that a given individual's disease was caused by the drug.

In addition to establishing the toxicity of the drug, the plaintiff also had to address the likelihood of other probable causes of her injury. In this regard, the court cites favorably its ruling in *Parker v Employers Mutual Liability Ins. Co. of Wis.*, in which it wrote:

There can be many possible 'causes,' indeed, an infinite number of circumstances can cause an injury. But a possible cause only becomes 'probable' when in the absence of other reasonable causal explanations it becomes more likely than not that the injury was a result of its action.

In cases like these, statistical evidence demonstrating both toxicity and causation, alongside trace evidence of the plaintiff's injuries, can suffice for findings of liability. This is unlike the Blue

²⁴ This is often referred to by the courts as establishing both 'general' and 'specific' causation. For more on this distinction, see Gold (2011).

²⁵ For an in-depth discussion of the legal issues surrounding differential etiology, see Sanders and Machal-Fulks (2001).

Bus case, in which statistical evidence of its market share, alongside trace evidence of the car accident, does not suffice for findings of liability.²⁶

In our view, one of the reasons why it is justified to treat toxic tort cases differently than the Blue Bus case is that the resulting incentive structure is significantly less problematic. A few contrasts between a standard toxic torts case and the blue bus case help make this clear. First, it is very rare that one can do better than statistical evidence (plus the profile of the ailment) in a toxic torts case. When buses hit cars there are often eyewitnesses. But there are no eyewitnesses to whether a certain chemical dumped by a company was in fact the cause of a particular case of cancer. This is especially clear given the latency periods of many chemicals: the etiology of a particular case of cancer can typically not be traced with assurance.²⁷

Just as the character of the dent contained no tell-tale signs of a particular bus company being the cause in the bus case, so the character of the ailment contains no tell-tale signs as to which of various possible causes are actual in a typical toxic tort case. In nearly every case there are possible causes of the cancer other than the chemical outputted by the relevant company and the best experts can do is make an informed judgment about the comparative likelihoods.²⁸ In the case of character evidence, Sanchirico's judgment is that the best incentive structure results from including character evidence at the *sentencing phase* but not the trial phase. An analogous strategy here would be to use it only at the phase of assessing appropriate punitive damages. But that strategy here would mean that punitive damages never occur, since there is no hope of getting a judgment of liability without relying heavily on comparative probabilities.²⁹ Moreover,

²⁶ We can dramatize the analogy with the blue bus case: The Blue Waste Company dumps blue molecules into the river of a town. There are also naturally occurring red molecules in the river. Both blue molecules and red molecules are known to cause mesothelioma, though the process of inducing mesothelioma takes many years and there is no method for directly determining the etiology in a particular case. The blue molecules outnumber the red molecules by a ratio of 5 to 1. It is incredibly unlikely that any case of mesothelioma is caused by the joint action of both red and blue molecules. A person from the town gets mesothelioma by drinking from the river. They sue the blue waste company for damages.

²⁷ For a discussion of these issues in the context of asbestos litigation, see White (2004).

²⁸ For an incisive (though opinionated) discussion of the ways in which courts can have difficulty parsing expert testimony on this point, see Gots (2005).

²⁹ We might note in passing that the suggestion of using market share at the "post trial" stage of a civil trial in the blue bus scenario would be particularly bizarre when combined with the practice of requiring that liability for one particular company be established at the "trial" phase. The suggestion would be that we ignore market share when assessing liability for a particular company but use market share to assess the amount of liability. As liabilities are typically not punitive but in amounts appropriate to damage done, this practice would make little sense. Moreover, while character evidence can be an appropriate guide to the malice of an act – "Was it done out of a stable disposition to wreak havoc or was it an out of character, one-off, aberration?" – market share can do nothing of the sort. Moreover, one crucial aspect of the role of using character evidence at either the trial or sentencing stage is this: The would be perpetrator reasons "If I commit bad act of type x, there is not only a chance that I will be caught with consequences y but moreover that conviction will then be used to magnify the potential penalties of future acts of type x". This kind of reasoning – which figures prominently in Sanchirico's discussion – has no natural analogue in the blue bus case. (As indicated earlier, it would be far better to have a practice where, in the absence of

the absence of anything like an eyewitness to the etiology is foreseeable in a toxic torts case, so if the fact of the ailment plus statistics is deemed insufficient, the company producing the chemical can foresee that no finding of liability against them will occur.

There are further contrasts. In the blue bus case, companies are disincentivized from having a majority market share – which is likely to be a bad thing from an economic point of view. Using statistical evidence in toxic torts is likely to disincentivize the dumping of chemicals, which is a good thing.³⁰ Meanwhile in the blue bus case, the amounts of money at stake would hardly justify the use of experts to assess the joint probabilistic impact of the relevant evidence, while in toxic tort cases, the amounts of money at stake provide ample justification for such a practice.

It is not surprising then that statistical evidence is allowed as the basis for findings of liability in toxic tort cases. As we have seen, there is excellent incentive-based reason to institute a system that allows statistical evidence to operate as a central basis for findings of liability in toxic torts even if such a system is a bad one in the blue bus case. So it should not be all that surprising that statistical evidence is treated differently in each context.

Some concerns remain however, even in the setting of toxic torts. There is the prospect of defendants avoiding liability on the grounds that, although their illegal activities clearly expose the plaintiff to a higher risk of injury, there are so many similarly situated defendants that identifying one is near impossible.³¹ Along similar lines, in cases where multiple companies are illegally dumping toxic waste it seems both unfair and unwise from an incentives point of view to allow a company doing 20 per cent of the dumping in an area to routinely pay nothing to compensate those who are injured.³² What such cases may ultimately show us is that, in toxic tort cases, environmental regulation and enforcement has a crucial role to play in the most effective and appropriate incentive structure.

evidence that incriminates one particular company, a trial phase established that the bus industry was responsible, with liability split within the local bus industry according to market share.)

³⁰ Unfortunately, even with these disincentives in place, pollution is often still economically rational for firms. For an in-depth empirical argument to this conclusion, see Shapira and Zingales (2017).

³¹ Arguably, one example of such a case is *Mullen v Armstrong World Industries, Inc.* (246 Cal. Rptr. 32 (Cal. App. 1. Dist. 1988)), in which 40 companies involved in the mining, manufacturing, and supplying of asbestos were sued by owners of homes containing their products. The court ruled that the plaintiffs could not bring suit against the companies because each produced or distributed a different asbestos-laden product with a different level of toxicity.

³² Concerns along these lines have motivated courts to find defendants liable on the basis of their market share of toxic or harmful substances (for instance, in *Sindell v Abbott Labs* 607 P.2d 924, 943 (Cal. 1980)). Such cases are distinct from the ones discussed above in that, in the cases discussed above, we focused on the question of whether it is more likely than not that the plaintiff's injuries were caused by a toxic substance. Once we establish this question, we then can ask the question at issue in *Sindell*, which is whether it is more likely than not that the defendant is liable for the toxic substance that caused the plaintiff's injuries.

When it comes to establishing liability for dumping by one or other company (as opposed to causes that have nothing to do with chemical dumping), the practice of relying on statistical evidence in combination of the trace evidence of the ailment seems eminently laudable from the point of view of incentive structures (and for what it's worth, we think it eminently laudable simpliciter). One lesson to be drawn is that philosophers should not be coming up with stories as to why it is generally bad to use statistical evidence as one's primary basis for conviction. A system – like the bizarre umbrella regimen – where trace evidence plays no role whatsoever will invariably create a bad incentive structure. The problem with that regime was not that it used statistical evidence, but rather that it failed to use trace evidence. Systems which use both statistical evidence and trace evidence (even if only non-individualized presence of a bad outcome) will not invariably create a bad incentive structure. And in cases where they create a laudable incentive structure, there may, at least in the context of civil cases, be little prospect of a strong countervailing argument against the practice. Perhaps philosophers have been hoping for too much by way of general arguments and diagnoses in this area.

Concluding remarks:

Let us sum up. Enoch et al. contend that there is a connection between instrumental considerations against the use of statistical evidence (due to perverse incentives produced by that use) and epistemological considerations dealing with sensitivity (due to judgments formed on the basis of sensitivity tests failing sensitivity tests): the counterfactuals that test for sensitivity are the counterfactuals whose truth is needed for law to provide proper incentives. We have told a very different story. The kind of plain conditionals that might naturally figure in ordinary reasoning relevant to the incentivizing effects of law are ones that *pass* sensitivity tests in many statistical evidence settings, including the canonical setting where market share is used to assign liability to bus companies when no eyewitnesses are around. Moreover the “damned if I do, damned if I don’t” reasoning that plays a starring role in their story will not be compelling in most cases where bad acts produce non-individualized trace evidence that would not be produced otherwise. There is nevertheless an incentive-based story against using market share evidence in the blue bus case (one that can be bolstered by various considerations of fairness), but it is not one that neatly aligns with the truth or falsity of various conditionals. It is instead one that can be probed by looking at the effects that different legal systems would have on certain conditional probabilities—the probabilities of legal repercussions conditional on agents’ various possible actions. Moreover, as we have seen in the case of toxic torts, instrumental considerations may break quite differently in different classes of cases, and in some of them it seems entirely appropriate to give statistical evidence a central role. Enoch et al.’s idea to give incentives an important explanatory role is a good one. But the connection between such explanations and sensitivity is extremely tenuous, and the resulting case against statistical evidence is quite limited.

References

- Black, B., & Lilienfeld, D. E. (1985). Epidemiologic proof in toxic tort litigation. *Fordham Law Review*, APA 7th ed. 52(5), 732-785.
- Blome-Tilman M. (2015), 'Sensitivity, Causality and statistical evidence in courts of law,' *Thought* 4: 102-112
- Cohen LJ (1977) *The Probable and The Provable*. Oxford University Press
- Dahlmann, C. (2020) 'Naked statistical evidence and incentives for lawful conduct,' *The International Journal of Evidence and Proof*.
- DiBello (2018) Trial by statistics. A high probability enough to convict? *Mind* 127: 1193-1218
- Dorr and Hawthorne. *Conditionals*, ms.
- Enoch, D., Spectre, L and Fisher T., 'Statistical Evidence, Sensitivity and the legal value of knowledge,' *Philosophy and Public Affairs* 40, 197-224.
- Enoch, D. and Fisher, T. 'Sense and 'sensitivity' -- epistemic and instrumental approaches to statistical evidence,' *Stanford Law Review* 67: 557-611
- Gardiner, G (2018) 'Legal Burdens of Proof and Statistical Evidence,' *The Routledge Handbook of Applied Epistemology*, Routledge, 179-195.
- Gold, Steve C. (2011) "The "Reshaping" of the False Negative Asymmetry in Toxic Tort Causation," *William Mitchell Law Review*: Vol. 37: Iss. 3, Article 6.
- Gots, R. E (2005). *From Symptoms to Liability: The Distinct Roles of Differential Diagnosis and Causation Assessment. For the Defense, DRI*. Available at <http://www.ictm.com/Downloads/dri-causation.pdf>.
- Hajek, A. (ms) 'Most Counterfactuals Are False'
- Hajek, A. (1989) 'Probabilities of Conditionals Revised,' *Journal of Philosophical Logic* 18: 423-428.
- Hajek, A. (forthcoming) 'Contra Counterfactualism,' *Synthese*,.
- Harris Jr, O. F. (1986), *Toxic Tort Litigation and the Causation Element: Is There Any Hope of Recognition*, 40 *Sw L.J.* 909.
- Lewis, D. (1973), *Counterfactuals*, Oxford University Press.
- Lewis, D. (1976) 'Probabilities of Conditionals and Conditional Probabilities,' *Philosophical Review* 85: 297-315.

Moss, S. (2018) Probabilistic Knowledge.

Nesson C. (1979) 'Reasonable Doubt and Permissive Inference -- the value of complexity,' Harvard Law Review 98: 1357-1392

Nunn GA (2015) 'The incompatibility of due process and naked statistical evidence,' Vanderbilt Law Review 68: 1407-1433

Pardo (2019) The Paradoxes of Legal Proof -- a critical guide. Boston University Law Review 99: 223-290

Posner (1999) 'An Economic Approach to the Law of Evidence,' Stanford Law Review 51: 1477-1546

Sanchirico CM (2001) 'Character Evidence and the object of trial,' Columbia Law Review 101, 1227-1311

Sanders, J and Machal-Fulks, J (2001). "The Admissibility of Differential Diagnosis Testimony to Prove Causation in Toxic Tort Cases." Law and Contemporary Problems. Col 64: 4, pps. 107-138.

Shapria, R. and Zingales, L. (2017). Is Pollution Value-Maximizing? The Dupont Case. National Bureau of Economic Research Working Paper 23866. Available at <http://www.nber.org/papers/w23866>.

Smith, M. (2017) 'When Does Evidence Suffice for Conviction?' Mind 127, 1193-1218.

Stalnaker (1976) 'Letter to Harper' in Foundations of Probability Theory, Statistical Inference and Statistical Theories of Science, Vol I, Dordrecht, 113-115.

Thomson JJ (1986), 'Liability and Individualized Evidence,' Law and Contemporary Problems 49: 199-219

Uviller, R (1982) 'Evidence of character to prove conduct -- illusion, illogic and injustice in the courtroom,' University of Pennsylvania Law Review 130: 845-891.

Wasserman D. (1991), 'The morality of statistical proof and the risk of mistaken liability.' Cardozo Law Review 13: 935-976.

White, M.J. (2004). Asbestos and the Future of Mass Torts, Journal of Economic Perspectives, 18:2, pps. 183-204.

Williamson, Timothy (2020) Suppose and Tell. Oxford University Press