

Statistical Evidence, Sensitivity, and the Legal Value of Knowledge

David Enoch, Talia Fisher, and Levi Spectre*

1. The Problem

A bus causes harm. In the first scenario, an eye-witness recognizes the bus as belonging to the Blue Bus Company. The witness, however, is rather imperfectly reliable; let's say that she's roughly 70% reliable in matters such as this one. The law has no qualms about accepting the eye-witness testimony as evidence, and indeed basing a positive finding that the bus was a Blue-Bus bus (and perhaps also that the Blue Bus Company is liable) on the testimony. In the second scenario, there is no eye-witness, but we have uncontested data regarding the distribution of buses in the relevant area; in particular, the Blue Bus Company owns roughly 70% of the buses there. Here, though, the law typically will not be willing to base a positive finding of fact – and certainly not liability – on just this kind of evidence, sometimes called statistical evidence. Indeed, in most jurisdictions it is not even clear that such evidence would be considered admissible or relevant.¹ And regardless of the reasons why (more on this shortly), it is an overwhelmingly common and strong intuition among practitioners and scholars alike that there is *something* suspicious about the second scenario, that there is *some* sense in which the market-share evidence is inferior to the eye-witness testimony. But of course, the case² has been devised so as to hold all apparently relevant features constant. The probabilities, in particular, are equal. This means that the chances of a finding of liability against Blue Bus Company being mistaken are similar in both cases. And after all, there is *something* statistical about the eye-witness testimony as well. She is fallible, of course, and out of an arbitrary set of,

* This paper was presented at the Uncertainty in Morality and Law Conference, sponsored by the Law and Philosophy forum of the Hebrew University in Jerusalem, in June 2011. We thank the participants for the valuable discussion. We especially thank our commentator, Amit Pundik, for very detailed, insightful, and helpful comments. Thanks to Cian Dorr, Ofra Magidor, Stewart Cohen and ... for helpful discussions. David Enoch's research was supported by the Israel Science Foundation (grant no. 136/09).

¹ See, for instance, Haw (2009) and the references there.

² For one discussion of this common hypothetical, see Redmayne (2008).

say, a hundred buses which she recognizes as Blue-Bus buses in similar circumstances, about thirty will belong to other companies. In what way, then, can we distinguish between the fact that the relevant bus was a member of a set 70% of which belonged to the Blue-Bus Company (not enough to ground a finding), and the fact that the bus was a member of a set 70% of which are accurately identified by the eye-witness (sufficient to ground finding)? How can we, then, accommodate in a theoretically respectable way the evidential distinction between the two scenarios we started with?³

The problem generalizes, of course. We can easily think of examples in criminal law or in the law of torts (or elsewhere), examples with different levels of probability (that are held constant in both scenarios), examples where there is or there isn't further evidence available, examples about the past or about the future (say, about dangerousness), and so on. And for some purposes, there may be important differences among these examples⁴. We will get back to more examples later on. For now, though, no further distinctions are needed. The problem of accommodating the distinction between statistical and individual evidence is a general one, and it seems to call for a general solution.

What should we say, then, about such cases? One possibility, of course, is to declare the law of evidence – and indeed, the powerful, uncompromising intuitions of pretty much all of its scholars⁵ – mistaken here. Perhaps, in other words, there is after all no convincing reason to distinguish between statistical and individual evidence. Though this is clearly a possibility⁶, it seems like it should be avoided if possible. We are, then, going to make an attempt – yet another attempt, perhaps, given the volume of literature here – to find after all a vindicating

³ One major issue we will not try to tackle here, at least not directly, is the reference class problem. We will assume, more or less throughout, that the statistical evidence latches on to the relevant frequencies. Such a simplifying assumption cannot be objected to in our context, as it arguably arises both for statistical and for individual, direct evidence (for instance, regarding the 70%-reliability of the eye-witness).

⁴ For a list of examples, and some insistence of the significance of the differences between them, see Redmayne (2008, 282-285)

⁵ See here Stein (2005, 78).

⁶ Schoeman (1987) is *very* sympathetic to this possibility, offering a debunking explanation of our intuitions to the contrary in terms of some cognitive biases. But we think that his pessimism about the distinction is premature.

explanation of this distinction. But how could this be so? Isn't it clear that in such pairs of cases, when we are holding all other things equal, there is no epistemic difference – no difference, in other words, when it comes to evidence, to justified or warranted or entitled belief, no truth-related difference, no difference that is knowledge-relevant – between the statistical-evidence and the individual-evidence scenarios? And isn't it clear that this is the only kind of difference evidence law should care about?

As things unfold, it will turn out that the answer to both these questions is "No". Even more interestingly, it will turn out that the underlying reasons for the two "no"s are intimately related. In the next section we distinguish between an epistemic and a practical way of addressing the initial challenge. On the former, there is after all an epistemic difference between statistical and individual evidence; perhaps, for instance, statistical evidence typically cannot – and individual evidence sometimes can – ground knowledge. On the practical strategy, though there is no epistemic difference of this kind (or perhaps even if there is no such difference), still the distinction evidence law draws here can be defended, because of some practical considerations, like perhaps some instrumental payoffs. In section 3, we leave the law of evidence behind, to discuss a general issue in epistemology – we start with one recent version of the lottery paradox, and we use it to motivate Sensitivity – roughly, the requirement that a belief be sensitive to the truth – as a necessary condition for knowledge. Armed with Sensitivity, we then return (in section 4) to statistical evidence, showing, first, the parallels between the general epistemological puzzle and the problem of statistical evidence, and second, that Sensitivity is a good start for solving the latter as well as the former. By the end of section 4, then, a (partial) solution to the problem of statistical evidence is presented, one that is in line with the epistemic strategy from section 2. Epistemology is complicated, though, and nothing about Sensitivity is uncontroversial or self-evident. In section 5, then, we give some more details, refining the sensitivity requirement and presenting a more precise version that is better able to accommodate some objections from the recent epistemological literature. We also show

how these details support the use to which we put Sensitivity in section 4. All of this may leave you wondering – even if we got all the details right – why the law in general, or evidence law in particular, should worry about such niceties. After rejecting one initially plausible answer ("Why, the law in general, and certainly evidence law in particular, should care about knowledge!"), we fall back on instrumental considerations here. That is the topic of section 6. But utilizing some recent work by Chris Sanchirico, we show that at least one family of instrumental considerations is – though distinct from the epistemic consideration captured by Sensitivity – still very closely related to it. That the epistemic and instrumental considerations coincide, then, is not merely a happy coincidence. In section 7 we discuss – in a somewhat preliminary way – the status of statistical evidence in morality rather than in law. This discussion – which is of independent interest, we think – is especially relevant here because it may be the beginning of an objection to our line regarding the legal status of statistical evidence, and in that section we address this objection. In section 8 we conclude by returning to the kind of example we started with, and indeed to some doctrines in the law of evidence. As we show in that section, our way of explaining the significance of the distinction between statistical and individual evidence gives the independently plausible results in some central cases mentioned in the literature. The discussion in section 8 is somewhat preliminary, and we hope to engage it – and doctrinal features and puzzles in the law of evidence – in more detail in future work.

Sometimes, a good way of appreciating the depth of a problem is by noting some flaws in the suggestions made by those thinking it is in fact much less profound. In this spirit – and also because doing so will help evaluate our own suggestion in following sections – we want to offer here a quick critical survey of some of the things that people – lawyers and theorists alike – sometime say in an attempt to vindicate the distinction between statistical and individual evidence⁷. The discussion will be quick and inconclusive – but such a discussion suffices, we

⁷ For a good survey that we found helpful here, see Ho (2008, 137-9). And for an earlier, much more critical discussion of many of these suggestions, see Schoeman (1987).

think, for its limited purposes here (namely, to help in appreciating both the problem and the advantages of our suggestion as a way of addressing it).

Sometimes, then, people mention other possibly relevant factors. Perhaps, for instance, the fact that there is in front of the court only statistical evidence is itself some evidence, perhaps evidence that no other evidence could be found; and perhaps that in itself is evidence against the plaintiff (or the prosecution)⁸. If so, statistical evidence should be accorded less weight, simply because cases in which it is presented tend to be cases in which the case of the party presenting it is weaker. This may be so – though much work would need to be done to show that this is so. But we can safely abstract away from all of this by insisting – as we did in the opening paragraph – that we hold all other things equal between the statistical-evidence scenario and the eye-witness one. Even after holding all else equal, still the intuition that there is an important difference survives. So explanations of the kind just mentioned are not what we are after.

It is sometimes insisted⁹ that there is an important difference between evidence that is genuinely about the relevant defendant (that the eye-witness said it was a blue bus is about the Blue Bus Company), and merely statistical evidence, that is thought to somehow be none of the defendant's concern (why is market-share in any way relevant to determining what happened in this specific case?). But there is, as far as we can see, no way of making sense of the about-relation that will vindicate this line of thought¹⁰. Because we are speaking of evidence, the only "about" needed here is the epistemological "about", the "about" of indication. And in both cases, as far as anything thus far said is concerned, the relevant piece of evidence does indicate that the bus was blue. In this sense, the statistical evidence too is about the Blue Bus Company. Now, there need be nothing objectionable about using such about-talk to capture the intuitive

⁸ Posner (1999, 1509).

⁹ Wright (1988, 1050)

¹⁰ Schoeman (1987, 183-4) gets close to this point, but not, we think, quite close enough.

distinction between statistical and individual evidence. It's just that doing so – without giving the kind of details about this about-relation that cannot, it seems, be given – amounts not to an explanation or vindication of the distinction, but rather just to giving it another name.

Some scholars talk in terms of the specific defendant, who is entitled not to be, say, punished for being a member of a reference class¹¹. Such talk is perhaps less convincing in the Blue Bus case with which we started but it does seem to capture something intuitive in other cases, such as the gatecrashers case¹², where it is uncontested that of, say, a thousand people attending a stadium event, only ten purchased tickets. If an individual – call him John – is sued, or even more clearly, if he is prosecuted – then finding against John merely on the strength of the (very strong!) statistical evidence here seems to be inappropriate. But though such a conviction would be inappropriate, it would be wrong to think of it as punishing John for being a member of a group many of whom crashed the stadium gates. In this case too, if we end up punishing John, we will be punishing him for crashing the stadium gates. It's just that we do not have a god-eye's view of the facts, and so we must determine – by relying on evidence – whether John did in fact crash the gates. And in doing this, the statistical evidence seems relevant – or anyway, if it is not, this remains to be shown. The point can again be made by noting that there is something statistical about individual evidence as well. Indeed, it is precisely here that it becomes tempting to insist – as some have¹³ – that really, at bottom, all evidence is statistical evidence. But presumably this doesn't show that in the eye-witness scenario we're punishing someone for being a member of the class of those recognized by the eye-witness – we punish them for committing the offense.

¹¹ See Colyvan et al (2001) (“Is It a Crime to Belong to a Reference Class?”), and Lempert (2001, 1669).

¹² David Kaye (1979, 104), Rhee (2007, 289)

¹³ .(..)Jonathan J. Koehler & Daniel N. Shaviro (1990, 263)

Relatedly, the point is sometimes made¹⁴ that the courts' primary obligation is to do justice in the specific case in front of them, and so that they are not entitled to sacrifice justice in the case in front of them just in order to achieve the result that is more efficient or is likely to minimize risk of error globally. But this doesn't help either, because the court which is determined to ignore all global effects and to attempt solely to do justice in the case, still has to use evidence – some evidence – to determine what doing justice in the case requires. And so far no reason has been given to believe that statistical evidence is any less kosher for this mission than individual evidence¹⁵.

Relatedly, it is sometimes suggested that relying on statistical evidence offends against the relevant person's autonomy, indeed perhaps even against her very free will and status as an agent¹⁶. By relying on statistical evidence to convict a gate-crasher, aren't we in effect saying that she was bound to crash, that it was not after all up to her, or some such? If so, aren't we – by relying on statistical evidence – in a sense degrading her? And isn't this reason enough not to rely on such evidence? There need be nothing wrong, we agree, with excluding degrading evidence, even when it is acknowledged as good epistemically, as genuinely probative, or some such. But this won't help here, because this line of thought does not plausibly generalize to all the relevant cases (think again of Blue Bus), and more importantly, because it incorporates a confusion similar to the one already highlighted above: The relevance of statistical evidence is merely as evidence; in a slogan, it's about epistemology, not about metaphysics. By taking something to be a reason to believe or to find that John crashed the gate we do not express any belief about it having always been bound that he would, or any such thing. We're just taking one

¹⁴ Lillquist, (2002, 140)

¹⁵ For similar reasons, talk of corporate punishment, or of collective punishment, or of the need to address the specific defendant rather than a group, will not help here. For many references, see Ho (2008, 139). Ho himself is guilty of similar mistakes, when he talks about relying on statistical evidence as intentionally taking a gamble at the defendant's expense; in cases of statistical evidence, "we saw an inadequacy in the evidence and we intentionally subjected the defendant to an open risk of injustice: we gamble on the facts at his expense." (2008, 142). But of course, there is always this inadequacy with fallible evidence, and so the only way we can avoid seeing this is if we are determined to ignore the obvious truth. And criminal procedure always involves intentional subjection of the defendant to a risk of injustice. Individual evidence is in no way better in this regard compared to statistical evidence.

¹⁶ Wasserman (1991-1992); Pundik (2008).

thing as an indicator of another. And surely, there's nothing special about that – we use statistical evidence about people's behavior (say, that most people stop at a red light) as data for our theoretical and practical reasoning (like deliberating whether or not to slow down as we approach an intersection) all the time. There seems to be nothing suspicious – let alone objectionable – about us doing so.

Sometimes there is also talk of what the public is or is not likely to accept¹⁷. Now, it is not clear to us what exactly the boundaries are of what would be acceptable by the public here (nor is it clear when – except in exceptionally high-profile criminal cases, perhaps – the public takes even the faintest interest in evidence law). But even assuming such doubts away, still we can safely bypass this consideration here. True, it is arguably important that the legal system enjoy some public trust (though questions may be raised as to the soundness of this as an aim independent of the legal system *meriting* public trust). And perhaps – though this is much more problematic – securing public trust can even sometimes justify catering to the prejudices of the masses. But if there is no other way of vindicating the traditional attitude towards statistical evidence, then this feature of public opinion is indeed a prejudice, and this at least renders suspicious the call to accommodate it. Furthermore, for our purposes here we can just assume the problem away by assuming that the public is going to have pretty much the right opinion about statistical evidence. In this (perhaps hypothetical) case, then, nothing about public opinion and trust can render justified an attitude towards statistical evidence that is not otherwise justified. Of course, if there is another justification for the traditional attitude towards statistical evidence, then (justified) public opinion may support it; but then, it will be primarily this other justification that does the relevant justificatory work. In this case too, then, we can safely ignore the public opinion argument.

Here is another proposal that initially may seem plausible, but must ultimately be rejected. If we were to bring before the court each and every person who came out of the gatecrasher

¹⁷ Nesson (1985, 1379).

stadium, and convict every one of them on the statistical evidence, we would be guaranteed to find guilty the 10 innocent people who purchased tickets. In non-statistical cases, though the probability of finding an innocent party guilty might be higher than in every one of the gatecrasher trials, we would have no such guarantee. Since such a guarantee is something we want to avoid, we accept only non-statistical evidence. But this line of thought cannot justify the full extent of the distinction, nor even to explain it. To see this, note just the following two points: First, in any criminal legal system, we are virtually (if not logically) guaranteed to convict innocents. The only way to avoid this result is to abolish a criminal justice system altogether. And it's hard to view the mere *logical* possibility of a criminal system not convicting and punishing innocents – a possibility whose probability is significantly lower than the probability that the next time you'll cross a road you won't make it to the other side – as making a serious normative difference. Second, think of a variant of the gatecrashers case, where for some reason we can only indict one person (perhaps all the others fled the stadium before the police arrived). In this case, relying on statistical evidence does not guarantee the conviction of an innocent person, but the intuitive reluctance to convict on the basis of just statistical evidence is fully present.

Let us again emphasize that this quick critical survey was meant to be neither conclusive nor comprehensive. Still, we hope it succeeds in giving a feel for the depth of the problem, and so also in motivating the search for a deep, interesting solution (and also, we hope, in giving you a reason to read another paper on the topic).

2. Two Kinds of Solution

Broadly speaking, we can distinguish two possible strategies of vindicating the distinction between statistical and individual evidence.

Instances of the first, epistemological strategy will engage in epistemological discussion, attempting to show that – at least in the relevantly paradigmatic cases – statistical evidence is

epistemically inferior to individual evidence. Such attempts can show, for instance, that statistical evidence never justifies belief; or (more plausibly) that it's harder for statistical evidence than it is for (probabilistically similar) individual evidence to do so; or that individual evidence can sometimes suffice – and statistical evidence cannot – for being entitled to hold on to a belief, or warranted in having a certain degree of confidence; or that statistical evidence cannot – or is much more unlikely to – render a belief rational; or that individual evidence can support knowledge, but statistical evidence cannot. All of these are arguably epistemic matters, as can be clearly seen when we distinguish between them and more practical matters. Thus, it seems rather uncontroversial that statistical evidence can render some actions rational; it can justify, for instance, certain gambles. (And when it comes to gambles, isn't it clear that it can justify them just as well as individual evidence?) But justifying actions or rendering them rational is one thing, and epistemically justifying beliefs is arguably another. The epistemic strategy insists that there is a difference between statistical and individual evidence that is of this latter, epistemic kind.

Of course, different attempts at this strategy may focus on different epistemic concepts (justification, entitlement, warrant, knowledge, rationality, epistemic reasons), and they may vary along other dimensions as well. But all of them have in common the insistence on the difference being roughly of this kind, rather than of the gamble-kind, or indeed of the instrumental kind we are about to get to.

Instances of the second, practical strategy accept – for the sake of argument, at least – that there is no epistemic difference between statistical evidence and individual evidence. Roughly speaking, as far as truth or conduciveness to truth is concerned, once we keep the probabilities constant across the two kinds of scenario, the game – this game – is over. But epistemology is one thing, and evidence law quite another. And so it is possible that there are practical reasons – for instance, instrumental reasons having to do with institutional features, with administrative costs, with differential incentives, and so on – why the law should take

individual evidence more seriously than it does statistical evidence. Now, practical justifications – certainly instrumental ones – are highly contingent. So it cannot be taken for granted that – having found a convincing instrumental story – it works in all and only the cases we intuitively think of as cases of suspicious statistical evidence. Much will depend here on the details. Given a sufficiently serious mismatch between the class of cases to which the relevant instrumental story applies and the class of cases we intuitively think of as cases of statistical evidence, we will be justified in rejecting the instrumental story as failing to capture the kind of general vindication we were after. On the other hand, given a sufficiently good instrumental story which doesn't apply in some peripheral cases of (what we pre-theoretically think of as) statistical evidence, we may declare them cases of direct, individual evidence after all.

With the distinction between epistemic and practical strategies¹⁸ (for vindicating the distinction between statistical and individual evidence) at hand, we can now generate the following prediction: If the best vindication of this distinction is along instrumental lines, it is likely to be law-specific. That is, if what justifies the differential legal treatment of statistical and individual evidence is essentially related to the instrumental payoffs of the law so treating it, then it is essentially tied to the law's so treating it. It is, after all, quite possible (and also rather plausible) that the instrumental considerations relevant to the law are different from those applicable to other institutions, or perhaps outside any institutional context at all. If, on the other hand, there is an epistemic vindication of the distinction, then it is likely to apply much more widely, indeed perhaps as widely as the relevant epistemic notion (justification, perhaps, or knowledge) reach. Going in the other direction now, if the problem arises much more widely

¹⁸ See Schoeman (1987, 187) for a similar distinction. Redmayne (2008, 245) also introduces this distinction, but he adds a third strategy, in terms of attacking the inference from the statistical evidence to the relevant finding. We fail to see how this forms a third kind of strategy here: Either the problem with the inference prevents it from establishing the relevant *belief* (on which the finding is based), or it does not; if it does, the problem is epistemic; if it does not, the problem seems practical; either way, the relevant cases fall into one or the other of the two strategies differentiated in the text. Notice that the practical strategy is not necessarily limited to instrumental considerations. Perhaps, for instance, the autonomy-line mentioned earlier in the text relies on non-instrumental, but still practical reasons that differentiate between statistical and individual evidence.

than in the law, an epistemic solution rather than an instrumental one seems to be called for. We return to this prediction in section 4.

3. Sensitivity: The Basic Idea

Forget evidence law for a second (or a section), then. Let's do some epistemology.

Think about the following version of the lottery paradox (for knowledge)¹⁹: In the first scenario, you buy a lottery ticket, where the chance of winning is (literally) one in a million. The winning ticket will be picked tomorrow. Tonight, do you know that your ticket will not win? The answer that seems overwhelmingly plausible to most is "No". You may know that it's highly unlikely that you will win; you may be justified in gambling against your ticket in rather high odds; but you do not know that you will not win (and this even if, as things turn out, your ticket does not end up winning).

Compare this, now, to the following, second scenario. In this scenario you buy a lottery ticket with somewhat better odds – one in ten thousand, perhaps. You hold on to it for a day. Now the winning ticket has already been picked, and you find the winning numbers in today's newspaper. Your ticket's numbers are not there. Newspapers are pretty reliable on such matters, but not, of course, infallible. Let's suppose that factoring in all the probabilistically relevant information here – the initial odds, the probability that the newspaper made a mistake, whatever else may be relevant – the probability that your ticket nevertheless won is now, oh, let's say one in a million. Do you now know that your ticket did not win? In this second scenario, the overwhelmingly plausible (and common) answer is "yes". (Indeed, it is hard to see how this answer can be avoided without deteriorating into a rather global kind of skepticism).

¹⁹ A more extensive presentation of the knowledge related lottery puzzles can be found in Hawthorne (2004). This section is survey-ish in nature – we do not pretend to be making an original contribution here. The view we are concerned with is one of several accounts that engage with various variants of Kyburg's (1961) lottery paradox. The literature on this subject is vast and we will not try to cover it here.

In the two lottery scenarios, then, we were careful to hold probabilities constant. Yet intuitively, at least when it comes to knowledge, there is an important difference between them – in the first scenario, where your evidence that the ticket hasn't won is just the odds of the lottery, you do not know that the ticket will not win. And in the second scenario, where your evidence partly consists of the newspaper item, you do know that the ticket has not won. Given that the probabilities are held constant – and indeed, we may stipulate further that you know the probabilities in both cases, so that known probabilities are also held constant²⁰ – what can possibly explain this difference?²¹

One plausible answer is in terms of some relevant counterfactuals²². What would you have believed, in both scenarios, if your ticket had in fact been the winning ticket? In the first scenario, you would have still believed that the ticket was not going to win. After all, your pessimistic belief was based on the statistical data that is still there, unchanged, even in the case where the pessimistic belief is false (because the ticket actually wins). In the second scenario, though, things are different. In that scenario, remember, you based your belief partly on what was written in the newspaper. And the newspaper – while in no way infallible – is still, so we're assuming, at least reasonably sensitive to the facts here. So had your ticket in fact been the winning one, in all likelihood this is what the newspaper would have said. And then, following the newspaper's lead, this is also what you would have believed. So it seems highly plausible to say that in the second, newspaper scenario, had the belief that your ticket won't win been false, you would not have believed it (rather, you would have believed that your ticket won,

²⁰ This further condition may be needed in order to alleviate some internalist worries (in the sense “internalism” has in epistemological contexts); we don't think we need to worry about them here. Some epistemologists, following Williamson (2000), (2009), may want to make a distinction here between epistemic probabilities and chances. Others, following Cohen (1988) may want to talk of a change in context due to salience of error possibilities for the knowledge ascriber (or in the epistemic subject's practical environment – Hawthorne (2004)). For our purposes here, though, we do not need to decide these issues.

²¹ You may think that time makes a difference here, as in the first scenario the relevant belief was about the future and in the second about the past. But this is not so. We can change the first scenario so that the winning ticket has already been picked, it's just that you haven't yet read the newspaper. In this scenario too you don't know that your ticket hasn't won.

²² Prominent proponents of this kind of view include Nozick (see below), DeRose (1995), Goldman (1975), and Dretske (1971).

as in fact it did). And so we have a distinction between the two lottery scenarios. There is a kind of counterfactual that differs in truth-value in the two cases: The counterfactual "Had the relevant belief been false, you would have not believed it" ends up being true in the second scenario, where knowledge is present, and false in the first scenario, where knowledge is absent.

This is not, of course, merely an interesting curiosity or a mere coincidence. For such counterfactuals seem to capture something that is intuitively of tremendous epistemic significance. Without committing ourselves to anything more precise at this point, we can say that when such a counterfactual is false – when, in other words, a true belief of yours is one you would have held on to even had it been false – then your belief (true though it may be) is not appropriately sensitive to the truth. Indeed, the fact that your belief is true may be thought of as a kind of a fluke – you, after all, would have believed it even had it been false. And so it may be thought that there's no genuine epistemic achievement here on your part – you just, as it were, lucked out. But where the counterfactual is true (where, in other words, had the belief been false, you wouldn't have believed it any longer), your belief does seem appropriately sensitive to the truth, you do seem entitled to some intellectual credit here, it is not a mere fluke that you believe truly (after all, had that proposition been false, you would have no longer believed it).

We can now introduce, then, Sensitivity:

Sensitivity (First Attempt): S's belief that p is sensitive =_{df} Had it not been the case that p, S would (most probably²³) not have believed that p.

And reflection on the two lottery scenarios lends initial intuitive support to the thought that Sensitivity is a necessary condition for knowledge; that non-sensitive beliefs do not constitute knowledge.

²³ For reasons having to do with counterfactual semantics, it will be useful to have this qualification in the official statement of Sensitivity. See (**) below.

More details are needed here, of course, and some will be supplied below (in section 5). What we have at this stage is just an intuitive case for an intuitive requirement – that beliefs be sensitive to the truth, if they are to count as knowledge. Regardless of the details – indeed, regardless of whether at the end of the day Sensitivity or something sufficiently close to it can be defended – it cannot be denied, we think, that Sensitivity captures something that is intuitively, pre-theoretically of considerable epistemic weight. And this lesson suffices for now, and allows us to return to the law of evidence.

4. Sensitivity and Statistical Evidence

Surely, by now you will have noticed the similarity. The two scenarios in Blue Bus (market-share-evidence, and eye-witness-testimony) parallel the two lottery scenarios (where the belief is based merely on the odds, and where it is also based on the newspaper, respectively)²⁴. The parallel is not just based on the intuitive similarity (the cases do have a similar feel to them, don't they?). For we are now in a position to say more. In both cases Sensitivity is a step in the right direction. We have already seen that this is so for the lottery cases in the previous section. Let's revisit, then, some of the examples of statistical evidence, this time armed with Sensitivity.

Suppose, then, that in both Blue Bus scenarios, we find for the plaintiff and against the Blue Bus Company. Where we do so based on the individual evidence – the eye-witness testimony – it seems like our finding is sensitive. Had it not been a Blue Bus bus, would we have

²⁴ The evidence-law literature on statistical evidence (or on the proof paradoxes) has recently come to appreciate this similarity with cases of the type of the lottery-case, but has not appreciated in full the significance of this similarity. Thus, Stein (2005, 67) mentions a lottery paradox in a related context, but deals with a version of the lottery paradox that is not relevant to our concerns; Redmayne (2008, 297) discusses our version of the paradox and explicitly draws the analogy between the evidence-law cases and the epistemological literature on the (relevant kind of the) lottery paradox, but he fails to notice the relevance of Sensitivity (rather, he discusses the related, but less appropriate here, Safety condition, and even that only in a very sketchy way); Ho (2008, 168-169) briefly mentions the similarity but fails to put it to theoretical use; and no one, as far as we are aware, discusses in this context Sensitivity in sufficient detail in a way that allows to vindicate the distinction between statistical and individual evidence, let alone to shed light on why the law should care about this distinction (as we do in section 6), or to show how this way of understanding the distinction can help shed light on some related doctrinal features (as we do in section 8). In the more philosophical literature on statistical evidence, the parallel has been made earlier and more often (see, for instance, Thomson (1986, 236-239) but there too without learning from it the right lessons. (And we discuss Thomson's view in some detail below, in section 6).

found the Blue Bus Company liable? Probably not. Our eye-witness is not infallible, of course, but she is pretty reliable, and so had it not been a Blue Bus bus, she would have probably not testified that it was; and in that case we would not have found the Blue Bus Company liable. So in this scenario, the finding is appropriately sensitive. 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% of the buses, it's just that the bus that causes the harm is no longer one of its buses (rather, it's 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 non-sensitive.

Or consider the gatecrashers case. Even if the percentage of gatecrashers among all those attending the stadium is quite high, still, if we convict John merely on the basis of the statistical evidence our conviction is not sensitive. For had John not been guilty of gatecrashing – had he been one of the small number of law-abiding ticket-purchasing people at the stadium – we would have still convicted him.

Let us return now to the two strategies for vindicating the distinction between statistical and individual evidence that were discussed in section 2. There we distinguished between epistemic and instrumental attempts at vindicating the distinction. We also noted there that if the phenomenon to be explained is wider than the law-of-evidence problem, this will count strongly in favor of the needed solution being of the epistemic kind. And we are now in a position to say the following: The problem – that of distinguishing between statistical and individual evidence even when the pieces of evidence are probabilistically on a par – is indeed much wider than merely the legal one. As can be seen from the lottery examples, the problem arises even where there is no clear institutional context of any kind, and where the instrumental considerations that

may apply are few, weak, and anyway different from the ones relevant to evidence law. But the problem is clearly the very same problem, and the reluctance to rely on statistical evidence is clearly the very same reluctance. This means that instrumental attempts at vindicating the distinction within the law of evidence – even if successful on their own terms – still fail to capture the full phenomenon to be explained, and so are not as good as explanations as may be hoped for. More positively now – seeing that the phenomenon is much wider, what we need is an epistemological vindication of the distinction between statistical and individual evidence. And focusing attention on Sensitivity does just that: For noticing that Sensitivity (or something very close to it) is a necessary condition for knowledge explains why there is something suspicious about statistical evidence across the board – in the legal context, in the lottery cases, and anywhere else where we care about knowledge²⁵ or about our beliefs being sensitive to the facts.

5. Some More Details

Enough has been said, we hope, to motivate Sensitivity, to show that it's rather plausibly considered necessary for knowledge, or at the very least (as will become clear in this section), due to some more elaborate necessary condition for knowledge, using it can neatly and easily distinguish cases of non-knowledge from possible cases of knowledge. Moreover, Sensitivity promises to explain – and vindicate – the distinction between statistical and individual evidence in an epistemological way, the way we want it vindicated given the scope of the phenomenon to be explained.

But more details are needed. In this section we elaborate on how the truth value of counterfactuals such as Sensitivity is determined; we also note some of the problems epistemologists have been having with Sensitivity as a necessary condition for knowledge, and also some of the intuitive advantages Sensitivity has – at least once we weaken the claim from

²⁵ Well, *do* we care about knowledge when we're doing evidence law? We discuss this question in detail in section 6.

requiring Sensitivity as necessary for knowledge to insisting that Sensitivity (or something in its vicinity) is an epistemically relevant condition; and we briefly discuss the relation between Sensitivity and an attempt to understand the distinction between individual and statistical evidence in explanatory terms.

Knowledge, Sensitivity, and (some) Counterfactual Semantics

Of the epistemologists²⁶ advocating Sensitivity as a necessary condition for knowledge, it is perhaps Robert Nozick (1981) who is the first to be explicit and detailed about what this requirement amounts to, and so for the most part we will follow his account.

Nozick's account of knowledge is influenced by the intuition that a belief that counts as knowledge should somehow track the truth, or the fact that the belief is about, an intuition that led several theorists (e.g. Alvin Goldman) to develop causal accounts of knowledge. However, Nozick was well aware of the difficulties any causal theory of evidence must face. We want to briefly mention some examples that nicely highlight some such difficulties, and to show how a Sensitivity-based account does better here (later on we will also show how variations on such cases calls for amendments to Sensitivity). Consider, then:

Fake Barn County:²⁷ Henry is traveling through fake barn county, a county filled with barn facades that from the road look exactly like real barns. Ignorant of the abundance of fake barns in this county, Henry looks at one of the few real barns and forms a belief that there is a barn yonder in the field nearby.

Although Henry's belief – there is a barn yonder in the field – is caused by a perception originating in a real barn, he does not, it seems, know that this proposition is true. After all, it is just by chance that he formed this belief while facing a real barn, for all he could tell he could have just as easily been facing a fake barn (a possibility that in the circumstances is not at all

²⁶ For instance, Carrier (1971); Goldman (1976); Dretske (1971); Nozick (1981: 167-288); DeRose (1995), (2010). DeRose does not pose Sensitivity as a necessary condition for knowledge, but he comes close enough (see note **).

²⁷ First published in Goldman (1976) - Carl Ginet is credited in the (1992) reprint of this paper.

improbable). Moreover, if he were facing a fake barn, he would have still formed the same belief. Fake Barn shows that a simple account of knowledge in terms of belief caused by the relevant fact (the one that is purportedly known) cannot be correct. But a Sensitivity-based account deals with Fake Barn with ease. For intuitively, had Henry's belief been false – that is, had he been standing in front of a barn façade – he would have still believed that there was a barn in front of him. His belief is not sensitive, and so does not constitute knowledge. And notice that in this respect Henry's case is different from the standard perception case, where someone – not in Fake Barn County – sees a barn; had this person not been standing in front of a barn, she would not have believed that she was, and so her belief is sensitive, and so may qualify as knowledge. But we may want more than just the intuitive determination about the truth value of a counterfactual. We need, that is, a semantic for counterfactual. And for our purposes here, the semantic Nozick is working with – probably the most influential suggestion for a semantic for counterfactuals in the literature – will do.

The intuitive idea is rather simple. A counterfactual conditional has a false antecedent (it is, after all, a counterfactual). So it tells a story about another way the world might be, or another possible world. Now, it is not enough, in order to falsify a counterfactual, that there are some possible worlds in which its antecedent is true and its consequent false. The counterfactual "Had my car not started, I would have come in late" is true, even though there are some possible worlds in which my car fails to start, and yet I make it in on time. It's just that in the possible worlds where not much is changed except for my car not starting, I come in late. Stalnaker (1968) and Lewis (1973)²⁸ precisify this intuitive suggestion. We order possible worlds in terms of their proximity to the actual world. And in order to evaluate the truth value of a counterfactual, we start out from the actual world and "travel" to the nearest worlds where the antecedent of the conditional is true. We then "check" to see whether or not the consequent is

²⁸ For a good presentation of Stalnaker's (and others') semantics of counterfactual conditionals, see Sider (2010) and Bennett (2003). The differences between Stalnaker's and Lewis's accounts will not concern us here.

true as well (in all those nearest worlds where the antecedent is true). If the consequent is true in those worlds, the counterfactual is true, if not, then it's false. A counterfactual is true if and only if in the closest possible worlds in which the antecedent is true, the consequent is true as well²⁹. Now, it would be great to have an account of the proximity relation – what determines how close one world is to another? – but let us concede here that despite there being some suggestions in the literature, we don't have such an account up our sleeve. Still, in many cases it is intuitively clear which of two worlds is closer to the possible world (in the intended sense of "closer"). Consider, for instance, the world in which my car doesn't start, and I end up taking a bus, arriving late on campus, to a world in which my car doesn't start, and a helicopter-mounted philanthropist grabs me and gets me to my office on time. It doesn't seem like a precise account of proximity is needed in order to determine that the latter world is farther than the former from the actual one, so that "Had my car not started, I would have come in late" is true.³⁰

Another advantage of Sensitivity is that it nicely deals with at least some Gettier cases: Think, for instance, of Gettier's famous Ford example: "Two other people are in my office and I am justified on the basis of much evidence in believing the first owns a Ford car; though he (now) does not, the second person (a stranger to me) owns one. I believe truly and justifiably

²⁹ A counterfactual with a metaphysically or logically impossible antecedent is treated as (vacuously) true on the standard Stalnaker-Lewis semantics.

³⁰ The counterfactual semantics utilized in the text is not without problems. Indeed, it has now been argued (by Alan Hájek, MS) that it follows from such a semantic – together with some facts about the chancy nature of the universe – that no counterfactual is ever determinately true, because for any consequent, and any closest possible world in which it's true, there is bound to be an equally close world in which it's false. Needless to say, we do not attempt in this paper to defend (let alone develop) a semantics for counterfactuals. The crucial thing for our purposes here is the truth value we intuitively tend to assign to the relevant counterfactuals. The more detailed semantic theory is used here merely as a heuristic. And regardless of the technical problems with this semantics, it seems to us that it's an adequacy constraint on any account of counterfactuals that it respect *some* distinction between (what we intuitively think of as) true and false counterfactuals.

We need not rely, for the present purposes, on determinably true as well as false counterfactuals. What we need is a distinction (on the basis of Sensitivity-type counterfactuals) between the two types of evidence. We can rely, then, on the difference between clearly false counterfactuals (statistical evidence) and counterfactuals that are not clearly false (individual evidence). In light of the account that we are offering below it will become clear, we hope, why this will do. A similar suggestion utilizes the parenthetical "most probably" in the original Sensitivity condition (as we've done above). The idea is that evaluating things from the actual world, one can say that probably a quantum blip (or another chancy occurrence) would not take place and so: had p not been the case, S would most probably not believe that p. This can also be made to cover worries having to do with mistakes of eye witness, newspapers, etcetera: "Had I won the lottery, the newspaper would most probably have reported that I won."

that someone (or other) in my office owns a Ford car, but I do not know someone does.” (Nozick 1981: 173) Nozick was correct to point out that my belief (in the Ford case) is not sensitive. Had there been no one in my office who owned a ford (suppose the stranger did not own one), I would have still believed that someone owned a ford.

Sensitivity, then, has much going for it. Yet Nozick was aware of counterexamples to Sensitivity of the following sort: Suppose that Grandma can just tell – when she sees her grandson – whether he's healthy. She just has a very good eye for such things. If some other conditions are in place (she's reliable in such things, perhaps she knows she is, etc.) it seems clear that she knows, when she sees her healthy grandson, that he's healthy. But now suppose that Grandma's son (Daddy) doesn't want her to get too worried, and so, when Grandson is sick, Daddy doesn't let Grandma see him, instead telling her believable stories about his excellent health. Still, when healthy Grandson comes over, Grandma knows, it seems, that he's healthy. And yet Sensitivity is not satisfied – had this belief been false (had Grandson been sick, in other words) Grandma would have still believed that he was healthy. Nozick responds by relativizing the counterfactual to methods of belief formation. True, in the case above Grandma would have believed that Grandson is healthy even had he been sick. But had he been sick, and had she still formed the belief about his health using the same method (that is, by looking), she would no longer have believed that he was healthy. And perhaps this is enough. We move, then (following Nozick), from Sensitivity to Sensitivity*:

Sensitivity*: S sensitively* believes that p =_{df} had p been false, and had S formed her belief about p by the same method she actually does, S would (most probably) not have believed that p .

Sensitivity* nicely captures the thought that the method by which the thinker forms the relevant belief has to be sensitive. It is not as clear, though, that Sensitivity* has all the other

intuitive advantages Sensitivity enjoys. For instance, it's not clear how well it does in terms of ruling out luck: Isn't Grandma just lucky, in other words, to be able to form the belief sensitively*, given that she would have just as easily formed the belief using another method (relying on Daddy's testimony), which is not sensitive? And the relativization to methods may entail other problems as well³¹. But for our purposes we do not need to comment further on this issue.

Sensitivity* – indeed, to a certain extent even Sensitivity – has undeniable advantages. It nicely deals with some lottery cases, it gets right some Gettier cases, it gets right some cases (such as Fake Barn) that a causal theory of knowledge gets wrong; and when it gets these cases right, it seems to do so for the right reasons. None of this means, of course, that Sensitivity* is indeed a necessary condition for knowledge. And in the next section we highlight some of the problems facing Sensitivity*, concluding that it probably isn't necessary for knowledge after all. But its many advantages should at the very least convince you that there's something right in the vicinity of Sensitivity*, that something like such counterfactuals captures something that is – even if not necessary for knowledge – still epistemically significant.

Problems for Sensitivity*

Its many advantages notwithstanding, Sensitivity – even Sensitivity* – is not at all popular among current epistemologists. One – perhaps the main – reason why is that Nozick's account of knowledge (of which Sensitivity* is a central part) violated the ever popular principle of epistemic closure, which reads: For all subjects S and propositions p and q, if S adequately infers q from p and S knows that p, S knows that q.³² The problem is that a belief that p may be sensitive (and so, given other conditions, qualify as knowledge), p may entail q (and S makes a correct inference from p to q thereby coming to believe that q), without the belief that q being

³¹ It requires from Nozick a criterion for the individuation of methods. And this means that Nozick's account is prima facie vulnerable to the generality problem that haunts reliabilists. See Conee and Feldman (1998). DeRose (1995, 20-21) proposes to leave Sensitivity as it stands and give method in belief formation a role in determining the closeness relation between worlds.

³² This formulation is based on Hawthorne (2004). There are many other variants.

sensitive, and so without it qualifying as knowledge (even when the other conditions needed for knowledge are in place). Now, it is controversial whether knowledge is closed under (known) entailment, and so whether it should count as a disadvantage of Nozick's account of knowledge that it violates closure. And when it comes to knowledge, much more needs to be said here, of course³³.

But as we already said, we are not going to insist that Sensitivity* is necessary for knowledge. Rather, we're going to insist that Sensitivity* is epistemically relevant³⁴ (and this will suffice for the use to which we put it in discussing statistical evidence). So it's important to see that objections from closure – while having at least some force in discussing knowledge – fail when it comes to the evidence-for relation. In other words, if S has evidence e for p (that does not entail p), and S correctly infers q from p, it does not follow that S thereby has evidence for q. Non-closure for evidence can be established probabilistically³⁵. But it may be helpful to give an intuitive example here³⁶: Your memory of parking your car in your driveway is evidence that your car is in your driveway; this memory, though, is not evidence that no one stole your car from the driveway; but that your car is in your driveway entails that no one stole it and drove it away; so

³³ Nozick thought it an advantage of his account that it violates the principle of closure and though the current trend in epistemology is to view non-closure as a reductio of any account that entails it, there are exceptions to this rule, e.g., Dretske (1970), (2005), Harman & Sherman (2004), Sharon and Spectre (forthcoming). Nevertheless, this much seems correct: without an independent reason to think that knowledge is not deductively closed, the intuitions motivating an account that entail non-closure had better be very robust (i.e. more robust than those that motivate closure accounts).

³⁴ DeRose (1995) and (2010) emphasizes similar claims. Though he does not think that Sensitivity is a necessary condition for knowledge, he claims that it is epistemically significant and is optimistic about a development of a more elaborate Sensitivity or InSensitivity condition that will further highlight its central role. As one of the prominent defenders of epistemic closure, he adds a contextual parameter to his InSensitivity-based knowledge account. Thus on this account an insensitive assertion may change the context of knowledge ascription so that "S does not know that p" is true in the ascriber's mouth (thus attempting to avoid non-closure).

³⁵ Assuming that *q* is evidence for *p* if and only if, the conditional probability of *p* given *q* is greater than the unconditional probability of *p* and less than 1 (i.e. $1 > \Pr(p|q) > \Pr(p)$), it is easy to verify that though *q* can be evidence for *p*, *q* is not evidence for a proposition that logically follows from *p*, e.g., $\neg(q \& \neg p)$. This is so because though $\Pr(p|q) > \Pr(p)$, the conditional probability of $\neg(q \& \neg p)$ given *q* is lower than its unconditional probability ($\Pr(\neg(q \& \neg p)|q) < \Pr(\neg(q \& \neg p))$). Proof (there are other variants in the literature since Carnap (1950)): (a) $\Pr((q \& \neg p)|q) = \Pr(q \& \neg p) \& q / \Pr(q) = \Pr(q \& \neg p) / \Pr(q) > \Pr(q \& \neg p)$ assuming $\Pr(q) < 1$; (b) From the Kolomogorov axioms it follows that $\Pr(A|B) < \Pr(A)$ iff $\Pr(\neg A|B) > \Pr(\neg A)$; and so (a) and (b) entail (c) $\Pr(\neg(q \& \neg p)|q) < \Pr(\neg(q \& \neg p))$.

In another paper, one of us has argued that the non-closure of the evidence-for relation can be established non-probabilistically as well – see Sharon and Spectre (forthcoming).

³⁶ Based on Vogel (1990).

the evidence-for relation is not deductively closed. And this means that even if Nozick's critics are right to reject Sensitivity* for knowledge (since it would entail closure-failure), they do not yet have even a prima facie argument against our use of Sensitivity*, that is, using it to pick out an epistemically significant property. More specifically, they do not have an argument against the idea that Sensitivity* picks out an important aspect of the *evidence for* relation.

There are other complications here that would have been more relevant had we insisted on Sensitivity* being a necessary condition for knowledge. Thus, Kripke devised some clever variants of the Fake Barn case, where even though Henry's belief that there's a barn in front of him is insensitive (for the reasons explained above), a perceptually-formed belief that there's a red barn in front of him *is* sensitive (perhaps because of a strictly enforced law against red barn façades, so that if there weren't a real red barn in front of Henry, he wouldn't have believed there's a red barn in front of him). And there may be variants of the Grandma case which are problematic for Nozick: For instance,³⁷ suppose that Daddy's technique to protect Grandma from the worries about her sick grandchild is not to tell her that he's fine, but rather to apply makeup to Grandson and bring him to Grandmother looking deceptively well. But now suppose that today Grandson is indeed well, and Grandma forms the belief that he is just by looking. In this scenario this belief is no longer sensitive*, but arguably it still seems intuitive that Grandma knows – when all goes well – that Grandson is healthy³⁸.

There may be other problems with Nozick's account of knowledge, of which Sensitivity* is a major part. But we think that we can afford not to say much more about these complications here, given the more limited use to which we put Sensitivity* in this paper, namely, the claim

³⁷ Williamson (2000) contains other examples and counter arguments to which DeRose (2010) is an attempt to reply on behalf of his (In)Sensitivity contextualist account.

³⁸ Some advocates of Sensitivity* (for knowledge) may insist, as suggested in conversation by Cian Dorr and Ofra Magidor, that in this case Grandma indeed does not know that Grandson is well. Indeed Grandma's vision is (wouldn't we say?) not to be relied on with regard to her Grandson's health. Nevertheless there are several issues here for the Sensitivity* advocate. One, we might wonder why we needed to amend the original Sensitivity condition and insist on Sensitivity*. After all, we could have said the same thing about Grandma's belief, i.e. that if Sensitivity fails it is not reliable. Two, it is not clear that the intuition here is not motivated by the theory. Further refinements of Sensitivity* seem to make the generality problem more pressing here. Having said this, we do not wish to settle the issue – if indeed there is a way to defend Sensitivity, all the better for our purposes.

that Sensitivity* captures an epistemically significant condition, that in many central cases a belief that satisfies Sensitivity* is epistemically better for it. None of this is meant to belittle the significance of such complications – perhaps, at the end of the day, they show that nothing like Sensitivity* (perhaps even no other counterfactual condition) can play a central role in an account of knowledge. Or perhaps they show that even if something like Sensitivity* does play a central role in an account of knowledge, it does so in virtue of some other, deeper story that renders it so relevant (one way of reading the next sub-section is along such lines). It's just that our ambitions here are more modest, and so we can afford not to say more here.

Sensitivity and Explanations

If a counterfactual is true, then (at least typically) there is a law-like connection between the antecedent and the consequent³⁹. And so a natural thing to say about the distinction between statistical and individual evidence – a point which is closely related to insisting on the counterfactual Sensitivity requirement – is that in the case of statistical evidence no such law-like connection obtains. Arguably, there is a law-like connection between the eye-witness testimony and the relevant truth, but not between the market-share evidence and the truth as regarding which bus causes the relevant harm.

A closely related point – indeed, given some plausible assumptions about the relations between explanations and counterfactuals, perhaps the very same point – can be put in explanatory terms. At least typically, individual evidence is explained by the truth of the fact for which it is supposed to be evidence. Thus, the explanation of why the eye-witness testified that it was a blue bus will – in standard cases – be at least partly in terms of the fact that it was a blue bus. Where this is not a part of the explanation of the testimony, we would take that very fact to undermine the testimony. The case of statistical evidence is different, though. Even when it was in fact a blue bus, this fact plays no role in explaining why it is that the Blue Bus Company

³⁹ For example, see Goodman (1947) or Nozick (1981, 690). Or see note ** below.

controls 70% of the relevant market. Similarly, in the first kind of lottery case, where the belief that my ticket won't win is based solely on the odds, this evidence is not explained by the fact that I won't win. But in the second case, where the belief that my ticket didn't win is based also on the newspaper report, this evidence is arguably explained partly by the fact that my ticket has not won. And so this observation generates the obvious hypothesis, according to which for a piece of evidence to be good, possibly knowledge-supporting evidence, it has to be explained by the fact for which it is taken as evidence.

A similar point may be made in the opposite direction. Suppose some evidence misleads you – that is, though E was evidence for p, it turned out that not-p. In the case of statistical evidence, such a case invites a “you-win-some-you-lose-some” kind of attitude. If 70% of the buses are owned by Blue Bus, and that was our reason for thinking that the involved bus was a Blue Bus bus, then we knew going in, as it were, that we were going to be mistaken roughly 30% of the time, and that's that⁴⁰. Tough luck this time. But when individual evidence misleads, the situation is different⁴¹. If we relied on the eye-witness in ruling against the Blue Bus company, and it turns out the bus that caused the harm actually belonged to the Red Bus Company, this discrepancy seems to call for explanation. Certainly, settling for a you-win-some-you-lose-some attitude seems out of place. In the terms Smith (2010, 16-17) introduces, this means that only individual evidence normically supports that which it is evidence for, so that when it misleads an explanation seems to be called for⁴².

⁴⁰ There are some probabilistic properties that are relevant to this distinction. As we cannot go into these matters here, we will peruse them (hopefully) in subsequent work.

⁴¹ A fuller discussion of mistakes than we need or can afford to conduct here would include - among other things - a discussion of the ways in which the mistakes of many procedures are not themselves entirely random (newspapers printing erroneous lottery results are not equally likely to print all erroneous results, etc.). Notice that such details - it seems reasonable to expect - will be closely connected to the requirement to explain the relevant kind of mistake, emphasized in the text.

⁴² Smith (2010) goes much further than merely noting the explanatory point in the text. Rather, he thinks we should think of normic support as grounding epistemic justification, indeed as doing so even against probabilities – so that one belief may be more justified than another even if the latter is more probable (for the thinker), so long as the former is better normically supported by the evidence. Relatedly, Mark Schroeder suggested to us in email correspondence that while statistical evidence can support *credences*, it cannot support all-out beliefs (which, it follows, do not supervene on credences). Furthermore, recent work by Timothy Williamson (2009) and (forthcoming)

A full discussion of the relevant epistemological issues may need to decide between some of these suggestions. Is normic support, for instance, what is ultimately needed for knowledge, or is it something like Sensitivity? Is what's doing the ultimate epistemological work here some kind of law-like connection, and if so, what kind exactly, and in which direction? In some contexts these are very important questions. But not, we think, in ours: For given some plausible hypotheses about the relations between law-like connections, the truth values of counterfactuals, and the nature of the relevant kind of explanation – all these different epistemological stories are sufficiently close for our purposes.

For our purposes it suffices that Sensitivity-like counterfactuals capture – often enough, in sufficiently central cases – an epistemically relevant feature of the distinction between statistical and individual evidence.⁴³ We do not claim explanatory ultimacy for the relevance of Sensitivity, and so even if what does the ultimate explanatory work here is something like the explanatory points mentioned in the previous two paragraphs, so long as something like Sensitivity is still epistemically relevant (and it is, given that both the sketched explanatory stories predict a difference in the truth values of the relevant counterfactuals for statistical and individual evidence), our vindication of the distinction between statistical and individual evidence goes through⁴⁴. At this point, this is partly a promissory note to be made good on by the discussion in

defending a sharp separation between chance and epistemic probability is relevant here (see note **). We do not need to discuss these interesting suggestions for our purposes here.

⁴³ In further work we hope to explore some related probabilistic properties that seem to fit well with the distinction between statistical and individual evidence. One such property was mentioned earlier (footnote **46) regarding the divergence in probability between possible mistakes, a property that seems to square well with the explanation suggestion.

⁴⁴ Nevertheless, let us make the following points about Smith's (2010) interesting suggestion. First, he explicitly addresses Sensitivity (...), rejecting it – if we understand him correctly – because misleading evidence that normically supports the relevant belief (and thus grounds epistemic justification) is not sensitive. But this, it seems to us, is beside the point: Of course *bad* or *misleading* evidence – individual evidence included – can fail sensitivity (if a bad or misleading eye witness testifies that p, it may be the case that she would have so testified even had it been the case that not-p. Indeed, this is one of the standard ways of discrediting eye-witnesses). The crucial point for us is that *even good statistical evidence* fails sensitivity. Second, Smith's suggestion attempts to explain epistemic justification using thoughts about what does and what does not call for explanation. And we agree – as stated in the text – that he's on to something important here, at least regarding the correlation between good, justification- and knowledge-grounding evidence and what mistakes call for explanation. Still, given the opacity of *calling for an explanation* – the question "What calls for explanation?" seems to us profound, and we do not know of eye-opening answers to it – it is hard to see Smith's contribution as explanatory progress. Talking about which mistakes call for explanation (rather than about which evidence supports which beliefs) does not seem to reduce mysteriousness. Perhaps the truth values of

the next sections – we hereby officially promise that nothing in what follows will depend on Sensitivity being the last epistemological word here. In fact, once this point is noticed, we can afford to use explanatory tests (Which mistake calls for explanation? Does the fact the evidence is for explain the evidence?) as proxies for the truth value of the relevant counterfactuals. The results of these explanatory tests and of applying Sensitivity directly go hand in hand sufficiently often to allow such a methodology (a useful one too, as will emerge below).

Should the Law Care about Sensitivity (or Knowledge)?

Let's recap. Using (one version of) the lottery paradox, we introduced and motivated the intuitive requirement that beliefs – if they are to count as knowledge – must be appropriately sensitive to the truth. We then formulated Sensitivity, according to which for A's belief that p to be sensitive is for it to be the case that had p been false, A would not have believed that p. And we suggested that Sensitivity is plausibly considered an epistemically relevant condition (even if not quite a necessary condition for knowledge, and even if there is some deeper-still epistemological story – perhaps in explanatory terms – explaining why it is that Sensitivity is relevant). We then returned to the topic of statistical evidence, presenting an epistemological vindication of the distinction between statistical and individual evidence, relying on Sensitivity; and we argued that given the lottery paradox and related contexts where the very same phenomenon – the reluctance to rely on merely statistical evidence – is present outside any legal setting, an epistemological vindication (rather than an instrumental one) is precisely the thing to look for. We then discussed in some detail some of the complexities and problems surrounding Sensitivity, and indicated how – at least for our purposes here – they are to be dealt with (or bypassed).

Sensitivity-like counterfactuals – or the law-like connections that support them – are better candidates for being the more basic explanatory story. Nozick (1981: 690) hints at the connection between his Sensitivity condition and Armstrong's reliability view of knowledge that involves law-like connections "between the belief that p and the states of affairs that makes it true."

But it is now time to address a remaining worry that may have been on your mind for a while. For even if all of this is right, you may wonder, and Sensitivity (or some version thereof) is indeed necessary for knowledge, why should the law of evidence care about knowledge? Why, in other words, should it make a legal difference whether a certain belief constitutes knowledge? In this section we first present the remaining worry in more detail (in section 6.1). It will prove convenient to postpone the presentation of our way of addressing it (in section 6.4) until after two interludes – the first (in section 6.2) is a brief discussion of Thomson's related suggestion regarding statistical evidence, and the second (in 6.3) is a presentation of Sanchirico's incentive-based discussion of character evidence. Our solution (in section 6.4) is going to concede – pace Thomson – that the law should not care about knowledge, or indeed about epistemology in general⁴⁵. And we are going to endorse a generalization of Sanchirico's instrumental reasoning. But we will show that the relation between his instrumental reasoning and the epistemological discussion from earlier sections is not a mere coincidence. The law should not care about knowledge, but it should care about what is relevant for knowledge – it should care about Sensitivity. Or so we are about to argue.

The Remaining Puzzle: Why Care about Knowledge?

It is important, of course, that courts not err too often. It may not be entirely uncontroversial how important this is; or which mistakes it is more important to avoid; or whether this is more or less important than some other important things. But no one doubts, we think, the importance of the courts' avoiding too many, too "big" mistakes. Whatever the functions of the law, whatever goods it can help achieve, its ability to do so depends on the courts not erring too often. And parties seem entitled to courts using procedures that will render mistakes that will hurt their (the

⁴⁵ For the most part we're going to be talking of knowledge, but with the point in the text understood as implicit: All we are officially claiming is that Sensitivity is relevant epistemically, that one can use it to make a distinction between two types of evidence, not that it's necessary for knowledge. So in claiming that the law should not care about knowledge, we will mean to say that the law should not care about the epistemic significance of Sensitivity, however exactly it is that Sensitivity is epistemically relevant.

parties') interests sufficiently improbable (of course, other considerations too may be relevant to determining the right procedures).

But statistical evidence can help improve the court's reliability. Indeed, it can serve to minimize error just as much as individual evidence can. In cases of the kind we were focusing on throughout this paper, the relevant piece of statistical evidence is probabilistically on a par with the relevant piece of individual evidence. Why exclude it, then? Is it really just because statistical evidence cannot ground knowledge or because on its basis one would believe the relevant proposition even if it were false? But why should the law of evidence care about knowledge? It should care, undoubtedly, about truth or the avoidance of error. But why is it important that courts base their findings on knowledge? More broadly, why should the law care about epistemology in general?

If you're not yet convinced that this is a genuine problem, let us make the following two points. First, perhaps (though we are not sure) there would have been some plausibility to the thought that the law should care about knowledge, with knowledge understood as some pretty basic, intuitively transparent notion. But we already know better. At least with the discussion of section 5 at hand (and certainly against the background of the literature it engages), we know that things are going to be very complicated here. And the thought that the law should care about knowledge thus understood – with all the complexities that are the main order of business of epistemologists but about which no one else cares (or, it seems, should care) – just loses any plausibility.

Second, it must be remembered that to insist that the law should after all care about knowledge is to be willing to pay a price in accuracy. Indeed, excluding statistical evidence amounts to excluding (what is often) good, genuinely probative evidence⁴⁶. And this means that

⁴⁶ Some empirical data may be relevant here. If it can be shown, for instance, that statistical evidence is typically – though not necessarily – less reliable than individual evidence, then the point in the text has to be changed accordingly. We do not know of any reason to believe that this (or any close) empirical speculation is true. And

the legal value of knowledge – if it has legal value, and if that value is what grounds the differential treatment of statistical and individual evidence – sometimes outweighs the value of accuracy; that, in other words, in order to make sure that courts base their ruling on knowledge we are willing to tolerate more mistakes than we otherwise would have to, and indeed a higher probability of mistake on this or that specific case. This just seems utterly implausible.

The problem here parallels one that has recently been receiving much attention in epistemology. For even in epistemology it is not clear why we should care about knowledge. There too, it seems, we should care about truth; and perhaps we should also care about the justificatory status of certain beliefs or inferences – whether, say, it's rational to have some belief given some evidence, or whether we are entitled to infer certain propositions from certain others, or some such. But we already know (at least since Gettier) that truth and justification do not suffice – not even together – for knowledge. So why should we care about whatever else is needed for knowledge? It makes sense, the thought goes, to aim at truth, and perhaps also at justification. But why aim at knowledge? This is, to repeat, a controversial question that has recently been receiving much epistemological attention⁴⁷. But notice that in our case the problem is much harder: For regardless of whether knowledge can be shown to have epistemological value, it is very hard to believe that it has legal value, indeed enough value to justify tolerating higher rates and probabilities of mistakes⁴⁸. The point applies equally to the

indeed, some empirical evidence – about the striking unreliability of eye-witnesses – pulls in the opposite direction. See ...

⁴⁷ See ...

⁴⁸ Notice that this remains so even if we engage in "knowledge-first" epistemology (see Williamson 2000), perhaps partly because of a (purported) constitutive relation between assertability and knowledge. This is why Ho's (...) way of addressing the proof paradoxes in the law of evidence seems entirely unsatisfying to us. Even given his Williamsonian assumptions, why should the law care about, say, assertability – indeed, why should the law care about it *enough* to tolerate a higher rate and probability of mistakes? Indeed for a Williamsonian there is a serious problem here having to do with reliance on evidence since the same body of evidence may diverge radically between probabilities as apposed to chances. For Williamson that there is no close world where a belief is false does *not* entail that there is a low chance that the belief is false. Thus one can know that p is true (i.e. the Williamsonian epistemic probability on a subject's evidence is 1) while the chance that p is true may be close to 0. Once this separation is imposed between chance and epistemic probability it is not clear (in or outside the courtroom) what one should do if, say, the chance that p is low and the epistemic probability is high (supposing it is crucial for an agent's practical deliberation whether or not p). For more on related issues see Hawthorne and Lasonen-Aarnio (2009), Williamson's reply (2009) and Sharon and Spectre (MS).

explanatory suggestions in section 5.3: Suppose, then, that statistical evidence cannot ground knowledge or even justification because mistakes based on it do not call for explanation. Why should the law especially care about avoiding mistakes that call for explanation? Mistakes that do not call for explanation seem – absent some story telling otherwise, at least – just as harmful to the relevant party, just as detrimental to the relevant social interests, etc. as mistakes that do call for explanation.

In this way, then, the story of Sensitivity as an epistemically relevant condition may be thought of not as a vindication of the distinction between statistical and individual evidence, but rather as merely a diagnosis of the relevant common intuitions, and indeed, perhaps even the beginning of a debunking explanation of these intuitions: This story helps to see what these intuitions track – something like evidence that can support knowledge; but now that we know that the law of evidence should not care about what these intuitions track, we should perhaps discard those intuitions, at least when it comes to the law. The Sensitivity-based epistemological story perhaps renders the relevant intuitions understandable, but not defensible as basis for legal policy.

We agree that knowledge is not something evidence law – or law more generally – should take an intrinsic interest in. Also, we see no way of supporting the (somewhat more plausible) claim that the law should take interest in knowledge (as opposed to reliably formed beliefs) instrumentally, because of some good effects it is likely to have. A different story is going to have to be told, then, if the distinction between statistical and individual evidence is to be vindicated. But that story, we will argue, is very closely related to the knowledge-story. In this way, though knowledge has no legal value, it will end up being indirectly relevant after all.

There is a general reason to suppose that a Safety theory will be unable to account for the difference between statistical and individual evidence in the law case. The reason for this is that a conviction will often depend on individual evidence that does not amount, nor does it come close to amounting to knowledge. Thus, just as in the statistical case, there will be close possible words (according to the safety theorist) where one believes falsely that the defendant is guilty.

Before establishing this, though, it will be useful, we think, to briefly discuss Thomson's related suggestion.

First Interlude: Thomson

In her influential discussion of statistical evidence⁴⁹ Judith Jarvis Thomson suggests that the difference between statistical and individual evidence should be understood causally. Individual evidence (such as eye-witness testimony) is causally linked in an appropriate way to the thing for which it is taken as evidence – in Blue Bus, it is the fact that the bus that caused the harm was blue that caused the eye-witness testimony, and (it seems) in an appropriate way. In the case of statistical evidence, though, no similar causal link is present. That the relevant bus was blue, for instance, in no way caused the market-share evidence. Thomson thinks⁵⁰ that such causal links with evidence are a necessary condition for knowledge. And she also thinks that they are necessary for justifiable legal fact-finding⁵¹ at least partly, it seems, because she believes that knowledge is a necessary condition for justifiable legal fact-finding (at least in criminal cases).

Now, we are about to criticize Thomson, and to highlight the differences between her view and ours. But let us not underestimate the ways in which our theory resembles Thomson's and draws on her contribution⁵². First, her attempt at a vindication of the distinction between statistical and individual evidence is – like ours – epistemological rather than instrumental. And hers too is partly motivated by analogies with the lottery paradox and the generality of the

⁴⁹ Thomson (1986).

⁵⁰ Thomson (1986, 230)

⁵¹ Thomson (1986, 244-245).

⁵² We also want to note here similarities between our Sensitivity-based account and Alex Stein's Principle of Maximal Individualization (PMI) (...). Stein's emphasis on counterfactualizability is naturally very close to our discussion of Sensitivity. Still, it's important to note here that the extent of the similarity depends on the details (both of Stein's PMI, and of Sensitivity, as developed in section 5); that Stein's account does not – and ours does – tie things here with the most general epistemological discussions; and that in Stein's case we can still ask why it is that the law should *care* about maximal individualization. We have a more elaborate story here, in the rest of this section.

phenomenon to be explained. Second, given some close relations between causation and counterfactuals, there is a similarity between Thomson's suggestion in terms of causation and ours in terms of some relevant counterfactuals⁵³. Third, as we are about to note below, there is a way of reading Thomson as primarily making a more general point, with the causal condition being only a particular instance of the general strategy Thomson is primarily advocating. If this is so, our theory can be seen as an attempt at another particular instance of the general strategy Thomson is after. Despite these similarities, though, we want to quickly mention some central problems with Thomson's theory. Spending some time on them will allow us to highlight some advantages of our own theory, and to motivate more clearly our solution (in section 6.4) to the why-care-about-knowledge challenge.

First, then, there's the problem that causal theories of knowledge have not been popular recently among epistemologists, and not, it seems, without reason. We will not discuss here the problems facing causal theories of knowledge (though see again the discussion of Fake Barn cases, above)⁵⁴. We just want to note, rather obviously, that if there are conclusive reasons to reject causal theories of knowledge in general, they also count against Thomson's employing such theories as a way of epistemically vindicating the distinction between statistical and individual evidence.

Second, partly for reasons related to some of the general objections to causal theories of knowledge and partly for other reasons, the causal apparatus seems ill-suited to capture the legal distinction between statistical and individual evidence. For instance, courts may sometimes need to accept evidence (expert witness testimony, say) regarding some mathematical truths. But it is very hard to see how the causal requirement can be met here,

⁵³ Given even more clearly close connections between causation and explanation, the similarity between Thomson's suggestion and the explanatory suggestions mentioned in section 5.3 is even more apparent.

⁵⁴ ...

given that mathematical truths are (arguably) causally inert⁵⁵. Also, causal links – even appropriate ones – can be notoriously complicated. And cases can easily be constructed – cases with multiple causes, independent causal chains, different facts that suffice causally only together but not each on its own, different facts each of which suffices causally alone, etc. – where it is not clear what follows from a causal theory, and to the extent that it is clear, the implications are intuitively unacceptable⁵⁶. Finally, consider the certainty case, where, for instance, no one at the stadium purchased a ticket. In that case, intuitively the evidence is still statistical (100% is yet another probability, isn't it?), but here it does seem sufficient for conviction. It is not clear, however, how a causal theory can accommodate this result: After all, there is no appropriate causal link between no one having purchased a ticket and John's gatecrashing. Thomson explicitly addresses (248) the certainty case, but if we understand her correctly, rather than showing how her theory can accommodate the desired result here, she proceeds by introducing an explicit exception for the certainty case. This, of course, is objectionably ad hoc. At the very least, a theory that would have the desired result in the certainty case as a natural particular instance (rather than as an ad hoc exception) would be better for it. And indeed, ours is a theory of this kind. A theory that utilizes counterfactuals rather than causal links is also immune, of course, to problems arising from the complexities of many causal links, and from the need to accommodate knowledge of causally inert stuff.

Third, whatever the plausibility of saying that it is morally wrong to punish someone unless we know they are guilty, surely nothing in the vicinity applies to torts or to other parts of private law. The burden of proof in such cases is typically said to be that of the balance of probabilities, and even if this talk is imprecise, so that a somewhat higher threshold is often in effect, still it is clear that the law is happy to base decisions in torts (for instance) on degrees of confidence that

⁵⁵ Causal theories of knowledge in general encounter difficulties in trying to accommodate mathematical knowledge. For the explicit claim that this spells doom for these theories rather than for mathematical knowledge, see Lewis (1986, 109).

⁵⁶ For such ways of objecting to causal theories of knowledge in general, see ... For such criticism specifically in our context – and indeed, specifically criticizing Thomson – see Pundik ...

are far too low for knowledge. And this doesn't seem like a morally outrageous situation either. So there's a general lesson here with regard to the legal-value-of-knowledge challenge: Clearly, the reluctance to rely on statistical evidence is the very same reluctance in criminal law and in other legal contexts. But this means that the reluctance cannot be adequately explained – not even in criminal law – by features that are unique to criminal law (like perhaps, on the suggestion we are currently considering, the requirement that we not punish someone unless we know they're guilty). Now, Thomson is well aware that the problem arises outside the criminal law, and at one point (245) she explicitly addresses the different standards required for proof. She insists that she does not require knowledge, but rather confidence over the relevant threshold that a guarantee is present, and she understands guarantees causally. But notice that on this suggestion, what does the work is not even knowledge, but rather the raw causal link itself – it is this link that is supposed to distinguish between legitimate individual evidence and illegitimate statistical evidence. But this strengthens, we now want to argue, the legal-value-of-knowledge worry.

In section 6.1 we claimed that knowledge has no legal value, that there is no reason for the law to concern itself with knowledge. We think that this is true whatever the right account of knowledge ends up being. But we think that Thomson's causal account makes this even more obvious. For why think that it should matter legally whether a certain piece of evidence is appropriately causally linked to the relevant fact? Clearly, the law should care about the reliability of the relevant piece of evidence. But holding reliability constant, isn't the insistence that the causal link should matter a kind of causation-fetishism? The worry that knowledge-fetishism was involved has been with us for a while now, perhaps, but causation-fetishism seems even worse, even more of a fetishism, than knowledge-fetishism. And this means that the move to "guarantees" (abstracting from knowledge) cannot help Thomson. In fact, it just highlights the most serious problem with her attempt at vindicating the distinction between statistical and individual evidence – it fetishizes causation.

More can be said here, of course, and some of what can be said can make Thomson's view somewhat more plausible. Often, for instance, Thomson emphasizes (e.g. 244) the intuition that knowledge – as well as legal findings – should be based in a not-merely lucky way on the evidence. And she seems to think of her causal theory as a way of fleshing out this anti-luck intuition⁵⁷. Though we reject (for the reasons specified above) Thomson's causal theory, we are happy to accept the (somewhat vague) anti-luck intuition. Indeed, our own theory in terms of Sensitivity can be seen as an alternative way of fleshing out this very intuition. If so, our theory can be seen as another, improved (we hope) attempt at the general line Thomson sketches here.

Second Interlude: Sanchirico

Criminal law has mixed feelings about character evidence. Such evidence is typically admitted in the sentencing stage, but usually not as evidence for conviction, and this despite the underlying suspicion that there too it can be, in a sense, good evidence, it can serve to make courts' decisions more accurate. Furthermore, this mixed attitude towards character evidence seems to most of us justified. It has proved very hard, though, to offer a convincing explanation of why this should be so.

In a fairly recent paper⁵⁸ Chris Sanchirico concedes (for the sake of argument, at least) that there's no way of vindicating this attitude towards character evidence if we think of the law of evidence as exclusively aimed at helping courts find the truth, or make factually accurate decisions. But Sanchirico suggests that we change the way we think of evidence law here. We should think of evidence law as also – perhaps primarily – being about supplying good incentives for primary behavior, behavior of agents outside the courts and the legal procedure more generally. And here, character evidence is problematic. This is so, because at the point

⁵⁷ Anti-luck epistemology is currently fashionable. See Pritchard (2005).

⁵⁸ Sanchirico (2001).

most relevant for incentives – when an agent is deliberating, say, about whether and how to break the law – his character as well as the relevant character evidence is already given. The character evidence that will be entered as evidence against (or for) him does not depend on his decision how to proceed right now. And this means that we have a problem. For ideally, in order to generate the efficient incentives here, we would want that person to know that the likelihood of his being (charged, and convicted, and) punished strongly depends on whether or not he decides to break the law here and now. The weaker the dependence, the less weighty the incentive supplied to him by the law not to engage in this specific criminal behavior. So admitting character evidence at the trial stage will be counterproductive in terms of incentives. And given some plausible assumptions about the difference between the trial stage and the sentencing stage, about which is more relevant for deterrence, and so on, perhaps this line of thought can begin to vindicate the above-mentioned mixed attitude towards character evidence.

Of course, a lot may be going on with character evidence. And it needn't be a part of Sanchirico's claim that giving the right incentives to primary behavior is the *only* normative consideration governing the rules regarding character evidence. But even if other considerations apply, still Sanchirico has succeeded in drawing attention to another kind of consideration, one that it would be foolish for a legal system to ignore.

Sanchirico's paper is not about statistical evidence, but about character evidence. There may be some similarities between the two, but we won't pursue them here⁵⁹. The important point for our purposes is that Sanchirico's general strategy can be easily applied to statistical evidence as well⁶⁰. Think, for instance, about John, a potential gatecrasher who is now deliberating, considering the options of purchasing a ticket, or perhaps gatecrashing, or perhaps

⁵⁹ In particular, character evidence may be thought of as a kind of intra-personal statistical evidence. And just as with statistical evidence there is an intuitive feeling that the evidence is not sufficiently directly about the relevant individual, with character evidence there is an intuitive feeling that the evidence is not sufficiently directly about the relevant specific action.

⁶⁰ A point overlooked, for instance, by Pundik's (*The Inevitable Efficiency of Using Racist Statistical Evidence in Court*) central line of argument.

going home and doing something else altogether. We are assuming, of course, that John has no influence on the behavior of the others at and near the stadium. This means that he has almost no influence on the relevant statistical evidence – the percentage of those attending the stadium who did not purchase a ticket is only to a miniscule degree influenced by the conclusion of John's deliberation. For all intents and purposes, he should think of it as already given. If so, though, our willingness to rely on statistical evidence almost entirely annihilates whatever incentive the substantive criminal law can give John not to break the law. For if the statistical evidence is strongly against him – say, because 98% of those attending are gatecrashers – John already knows that he will be convicted, regardless of whether or not he buys a ticket. And if the statistical evidence is not strongly against him, he knows that it will constitute strong exonerating evidence, whether or not he is guilty of gatecrashing. Either way, then, he might as well go ahead and gatecrash – whether he does or not will have very small influence on his chance of being punished.

Of course, things are a little more complicated than that. For one thing, even if we accept statistical evidence and are willing to rely on it unqualifiedly, still there is always also the possibility of individual evidence becoming available as well. And this kind of evidence may have better effect in terms of incentives (think, for instance, about the possibility that John goes home, and then has an alibi; or that he purchases a ticket, and keeps it for proof; or that he gatecrashes, and is videotaped climbing the fence). Furthermore, in some cases incentives work in a somewhat more complicated way – think here about Blue Bus Company, where presumably the important incentives have less to do with the deliberation of a specific agent wondering whether to break the law, and more with economic deliberation of organizations, questions about what precautions to take, for how long to let the drivers drive, what level of activity is optimal, etc. And it is an open question – not one that can be determined in a general way and a priori – how the different considerations interact in specific real-life cases (we return to this below). But none of this undermines the modest way of generalizing from Sanchirico's

lesson about character-evidence: At least one important normative consideration governing the advisability of relying on statistical evidence is the fact that relying on it will render the primary-behavior incentives the law gives less efficient and accurate than they would otherwise be. And this, of course, counts against relying on statistical evidence. And because there is no similar incentive-corrupting effect to relying on individual evidence – even individual evidence that is probabilistically indistinguishable from the relevant piece of statistical evidence – we have here an instrumental vindication of the distinction we were out to vindicate.

Solution: The Instrumental Significance of Being Sensitive

And so at this stage we find ourselves in the following predicament: The initial phenomenon to be explained – the reluctance to rely on statistical evidence – is broader in scope than just in the case of the law of evidence, and it applies even in more purely epistemological settings (where nothing like the instrumental considerations applying to the law is relevant). An epistemological explanation is thus called for, and we tried to give one in terms of Sensitivity. But none of this seems to be the kind of stuff that should matter to the law, certainly not in a way that could justify tolerating a higher rate of inaccuracy. Here what is needed, it seems, is an instrumental account, one having to do with incentives, and we suggested one such story (following Sanchirico on character evidence). But of course, nothing like this instrumental story can help with the lottery paradox or other non-legal cases where talk of incentives seems out of place (or at least, if something like incentives can still be relevant, it's going to be relevant in a very different way). Are we stuck, then? Furthermore, is it mere coincidence that the epistemological and the instrumental considerations align so neatly, at least when it comes to the law?

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 don't, they won't.". 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 conditionals⁶¹. Suppose that John proceeds to crash the gates. Then his conditional "If I don't crash the gates, they won't 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, still both of them stem from the same source – Sensitivity-style counterfactuals. These are needed both for knowledge (or are in some other closely related way epistemically relevant), and for a reasonably efficient incentives-structure. While neither the epistemological story nor the instrumental one depend on each other, they are not totally independent of each other either – both of them depend on Sensitivity and related counterfactuals⁶².

And so what we end up with is the following rather complicated story. There is a need for an epistemological story, one that will treat lottery cases and legal cases (and other cases too, of course) alike. Sensitivity and its epistemic significance does that. There is also a need for a practical, most probably instrumental story, one that will vindicate the legal significance of the distinction between statistical and individual evidence without resorting to knowledge-fetishism. The generalization of Sanchirico's account does that⁶³. But that account too relies on the truth of relevant counterfactuals, indeed the very same counterfactuals the epistemological account relies on. Sensitivity is (a part of) the answer to both the epistemological and the practical questions.

⁶¹ Subjunctive conditionals are usually considered as entailing their counterpart material conditionals, but not the other way round. A classic example: "if Oswald didn't shoot Kennedy, someone else did," is true, but "had Oswald not shot Kennedy, someone else would have", is (arguably?) false. See Adams (1970).

⁶² As is often the case with explanations of coincidences, one may still ask the question whether the explanans itself is a mere coincidence. Is it, in other words, mere coincidence that Sensitivity and related counterfactuals are relevant both practically and epistemically in this way? Or is there perhaps some even-deeper story that can be told here? We do not know, but we can't deny that it would be *especially* nice if such a deeper story were to be found.

⁶³ Our focus on Sanchirico's account neither entails nor presupposes, of course, that no other considerations can contribute here. But for any other account it will have to be checked whether it coincides with the epistemic story of Sensitivity as Sanchirico's does.

Notice that on this account, when it comes to policy recommendations what does the work is the incentive-story, not the epistemological one (otherwise, we really would have a case of knowledge-fetishism). If there are cases, then, where the instrumental payoffs the incentive-account relies on are not in place, or if they are in place but are outweighed by other instrumental considerations, then even if relying on the relevant piece of evidence would violate Sensitivity, we do not see a practical reason not to rely on it. And the extent of the overlap between the epistemological considerations and the instrumental ones is to a large extent contingent. Perhaps, if the overlap is significant enough, there are second-order considerations (having to do with administrative costs, perhaps, or the instrumental value of the simplicity of the relevant legal rules) against relying on (non-sensitive) statistical evidence even in cases where other instrumental considerations do not suggest so. But it is quite possible that instrumental considerations will sometimes just not be there to back up the epistemological ones to the degree necessary to compensate for the loss in accuracy that is always involved in ruling out probabilistically respectable evidence. We think of this implication of our view – that it concedes that, and explains when, taking statistical evidence seriously and perhaps also legal reform may after all be necessary – as an advantage.

Before returning to the law of evidence, let us discuss an important objection.

Statistical Evidence for Resentment? An Objection

Perhaps, if your spouse is cheating on you, and if you have sufficient evidence that this is so, you are justified in resenting him or her for that. Suppose now that the evidence you have is just very strong statistical evidence: Say, the very high percentage of unfaithful spouses among the relevant social milieu. It would not be appropriate for you to resent your spouse for cheating on you based solely on such statistical evidence. Resentment that is exclusively based on such statistical evidence is – we take it as a robust intuition – inappropriate.

But for us, this is the beginning of an objection. Here's why: In the legal case, we argued, what should determine legal policy is not epistemological but practical considerations, in particular considerations having to do with incentives. But when it comes to whether resentment is appropriate, incentives seem irrelevant: The question we are asking now is whether merely-statistically-supported resentment is justified, not whether some action (like punishing, or *expressing* resentment) can be justified in those circumstances. So it is hard to see how incentives can be at all relevant. And so, if the only vindication we could offer of the reluctance to rely on statistical evidence is in instrumental terms, and if no instrumental story applies to the pure case of resentment alone (with no relevant action), we seem to be committed to the conclusion that nothing justifies the reluctance to rely on statistical evidence when it comes to resentment. And this conclusion contradicts what we characterized in the last paragraph as a robust intuition. Hence the objection.

The same point can also be put in the following way: The reluctance to rely on statistical evidence in the resentment case seems to be the very same reluctance as the reluctance to use statistical evidence in Blue Bus, Gatecrashers, and the others. So we are looking for a vindication that will apply to all of these⁶⁴. But our suggested vindication doesn't apply to the resentment case. So it's not an adequate vindication⁶⁵.

This objection raises the much wider issue of the appropriate evidence law (as it were) for morality. Thus, we can ask, for instance, what the burden of proof is that is required for appropriate resentment – beyond a reasonable doubt? Moral certainty? Knowledge⁶⁶?

⁶⁴ Indeed, at this point you may want to re-evaluate some of the alternative suggestions regarding statistical evidence. In particular, you may want to reconsider autonomy-based explanations, for they seem to get the resentment cases well (see here Pundik ...). But let us remind you that they face their own problems – most importantly, that of conflating metaphysics and epistemology (emphasized in section 1 above), and the fact that they can't accommodate the full scope of the phenomena themselves – namely, they seem utterly irrelevant to purely epistemological cases, like lottery cases.

⁶⁵ We thank Amit Pundik for pressing this objection on us in an especially forceful way.

⁶⁶ For instance, is resentment based on justified, true but Gettierized belief necessarily inappropriate? Can resentment be appropriate in cases where knowledge is absent for roughly the reasons – whatever exactly they are – for which knowledge is absent in the Fake Barn case? For some discussion of related questions (though in the legal, not the moral, context), see Pardo ...

Epistemic justification? Something else? Similarly, we can ask whether when it comes to morality there is anything equivalent to inadmissible evidence – in some jurisdictions, for instance, evidence illegally obtained is inadmissible in a criminal proceeding⁶⁷; but is there something inappropriate about merely blaming (or indeed praising) someone for an action, where our evidence that she has indeed performed the relevant action (or our evidence regarding the nature of the action performed) was immorally obtained? Questions about the status of statistical evidence when it comes to morality, as in the paradigmatic example of evidence relevant for resenting, are also among these questions of the evidence law of morality⁶⁸.

Quite astonishingly, we do not know of any systematic discussion of these questions (we are not even aware of a less-than-systematic discussion of them). And we cannot here engage these issues in a satisfactory way⁶⁹. So we are going to have to settle here for noting the following three related points. Let us emphasize that in making these points our ambition is rather limited – we are merely attempting to show that the objection we opened this section with to our vindication of the suspicion towards statistical evidence is far from conclusive. At the end of the day, it may exert a price – it may cost us some plausibility-points. But – as the following paragraphs show – the price in plausibility points is quite manageable.

Pardo introduces Fake Cab cases – an analogue of Fake Barn cases, where a legal finding is based on the testimony of an eye-witness – or by a jury – who is epistemically in the same spot as Henry in Fake Barn. Pardo presses the powerful intuition that such a finding would be wrong, and we tend to agree. Does this mean that here knowledge should after all matter to the law? We think that the answer is still “no”, and for reasons our account nicely explains. Something like Sanchirico’s incentive-based story applies. Once again, then, the epistemological story is not itself something the law should care about, but it goes hand in hand with an incentive story that is.

⁶⁷ ...

⁶⁸ Here is another relevant kind of question. When it comes to punishment, some discussion has been going on about whether – given suitable epistemic conditions – there is something necessarily wrong with “pre-punishment”, punishment that temporally precedes the act that calls for the relevant punishment. See, New (1992), Smilansky (1994), Statman (1997), and Sorensen (2006). How about pre-*resentment*, though? If you know that someone is about to wrong you, is it appropriate for you to resent them before they actually wrong you, simply because (as you know) they will? This question seems to us independently interesting (and the discussions of pre-punishment do not, as far as we know, address it.) But it is also relevant here: For the intuitive objection to statistically-based resentment seems to distinguish between statistically-based pre-resentment (more objectionable) and statistically-based post-resentment (objectionable, but not quite as much). A full discussion of morality’s evidence law would have to address this asymmetry as well.

⁶⁹ At least one of us hopes to do so in future work.

First, then: We insisted that the law should not care about knowledge, or Sensitivity, or which errors call for explanations. In less metaphorical terms – epistemological differences should not all by themselves make a legal difference. What the suspecting spouse example shows (perhaps) is that morality does care about epistemology, that it does make a moral difference whether one’s belief qualifies as knowledge, or is sensitive, or some such. But this doesn’t show that the law should also care about epistemology. Perhaps epistemology is something morality should and the law should not care about. In other words – those used three paragraphs ago – perhaps the reluctance to rely on statistical evidence in the legal and the purely moral case is not after all quite the same reluctance. We are not sure how much weight we want to place on this response, though, because of some plausible connections between law (especially, perhaps, criminal law) and morality. Perhaps, for instance, if punishing cannot be justified unless blaming and resentment are justified, so long as morality cares about knowledge (or some such), so should the law. But the details of such a story would have to be fleshed out and defended. At the very least, then, more work needs to be done here if we are to have a full objection to our theses.

Second, some examples of cases where resorting on purely statistical grounds is clearly inappropriate can be dealt with not by rejecting relying on statistics, but rather by insisting on doing it well. This is so, for instance, in the suspecting spouse example. For surely, within a close relationship, each spouse has much, much more information on his or her spouse than the information “he’s a married man” or “she’s a married woman”. Ignoring all that other, individual evidence about one’s spouse and relying exclusively on the statistical evidence is an instance of a statistical fallacy. The more specific, individual data arguably probabilistically screens off⁷⁰ the most general, statistical information⁷⁰, and so, given that the suspecting spouse is relying (as she should) on the fuller evidence, she shouldn’t also rely on the statistical

⁷⁰ For more on screening off, see, for instance Kelly (2005, 26-7) and the references there **or find what he refers to, etc.**

evidence⁷¹. But we do not want to create the false impression that this consideration – conclusive as it may be in diffusing the suspecting spouse example – can diffuse the objection as a whole. This is so because cases can be constructed when no such screening-off individual evidence is present, and in those cases too resenting on purely statistical grounds seems objectionable. We need to say more, then.

For our third point here, then, think again about the suspecting spouse case, and for now put to one side worries about which evidence screens off what other evidence. In that case, there is a close personal relationship in the background, one that is normatively rich – it includes many duties, commitments, entitlements and so on that are either a constitutive part of that relation or that are grounded in it. But then, it is very natural to read into this rich normative structure also the commitment of one spouse not to resent the other on purely statistical grounds, and indeed an entitlement of each spouse not to be resented by the other on such grounds. Personal relationships are, well, personal, and it is plausible that they include also a commitment to relying (only, or anyway mostly) on personal, specific, individual evidence in one's interaction with the person one is in a relationship with.

Well then, can't we improve the example so that no personal relationship of the kind that could ground such a commitment is present? The answer, it seems to us, is that this cannot be done, because resentment itself is personal, and is arguably impossible in a perfectly impersonal setting. The *judgment* that someone has acted wrongly is, of course, possible even without any personal relationship. But the richer response we call *resentment* seems to presuppose a personal relationship (perhaps sufficiently broadly understood). And the colder,

⁷¹ Notice that this is *not* just an instance of the requirement to base one's beliefs on the maximal evidence one possesses – that would be consistent with relying on the individual as well as the statistical evidence. The point is, rather, that once one relies (as one should) on the individual evidence, relying on the screened-off statistical evidence would amount to double-counting.

This general point may be reflected in the tendency of courts to rely on the most individualized evidence available. See: Fischer (1999, 1137).

less personal, less engaged judgment that someone has acted wrongly is possible and is not (as) intuitively objectionable even when based on purely statistical evidence.

Thus, we can accommodate the robust intuition – that resentment is inappropriate if based solely on statistical evidence – without compromising the explanatory story we put forward throughout this paper. In fact, the points just made can be seen as a way of rendering plausible the suggestion above, according to which the law should not care about epistemology but morality should – morality should, at least when it comes to attitudes like resentment, because resentment is personal, in the way that the law is not.

Still, as noted above, we may be paying here some price in plausibility points. This is so, first, because substantive assumptions have been made (with very little argument supporting them here) about resentment; second, because some details may be doubted here (perhaps, for instance, someone may want to argue that the law, or at least the criminal law, is after all personal in the relevant ways); and third, because examples may be constructed where the line suggested here may not be clearly plausible. For instance, if you read in the newspaper about serious wrongdoing in a distant village, can you resent the wrongdoers? The answer is not clear, but if it is "yes", then this may be thought of as some reason to doubt that a personal relationship is needed for resentment. But such examples – even if not clearly compatible with what's been said – are not *clearly* incompatible with that either, and so the loss in plausibility points – if there is such – is not that troubling, given our theory's other advantages (in particular, its success in both giving a unified epistemological account for lottery cases and legal cases like Blue Bus, and in accommodating the fact that the law should not care about epistemology).

Applying It All to Evidence Law Doctrine – a Sketch

This concludes, then, our vindication of the suspicion with which the law views statistical evidence, and in particular, the distinction it often draws between statistical and individual evidence. But much more can be said. We could now survey relevant doctrines in the law of

evidence, show how the story above fares in explaining them, and perhaps also where the story above calls for reform. But we cannot conduct such a detailed discussion here – we intend to do so in a separate paper. Instead, let us very briefly mention just some examples along these lines.

The Basic Examples, and Some Complications

We have already told our story regarding the basic examples. In the statistical-evidence-scenario of Blue Bus, the fact that we would have found against the Blue Bus Company even had it not been a Blue Bus bus renders our finding epistemically problematic, perhaps not a case of knowledge, perhaps not even a case of justified belief. This, though, is not something the law should care about. What the law should care about is the fact that if it bases a finding solely on the basis of the market-share evidence, it mis-shapes the incentive structure of the Blue Bus Company (and indeed, of the Red Bus Company as well). And this is not so in the eye-witness case (where the probability of a wrong decision is just as high as with the market-share evidence). This is why the law should discriminate between the statistical and the individual evidence in this case. Similarly – *mutatis mutandis* – for the Gatecrashers case.

But there are many more options than this way of putting things suggests, and many more complications have to be taken into account in shaping and evaluating evidence law doctrines. For one thing, even if we've succeeded in showing why the law should take statistical evidence less seriously than it does individual evidence, we haven't said anything about *how much* less seriously. The answer is likely to depend – among other things – on the relative value (in the circumstances) of avoiding the incentive-problem we highlighted on one side, and achieving better accuracy on the other. And these values may in turn depend on further factors, like perhaps the degree of probabilistic reliability of the relevant statistical evidence: The more reliable it is, the heavier the price in terms of accuracy that will have to be paid if the evidence is excluded (or even downgraded). In this way, the treatment of statistical evidence recommended

by our analysis may differ widely across cases, with such factors as: statistical evidence substantiating different levels of probability (as just noted); statistical evidence relating to past or future events; Cold hit statistical evidence (that is, self-standing, "out-of-the-blue" statistical evidence) versus confirmatory evidence (that is, cases where we already have a suspect, and the statistical evidence is used just to confirm the suspicion); statistical evidence relating to different strands and dimensions of liability (causality versus negligence; civil liability versus criminal culpability etc.); statistical evidence relating to different phases of trial (liability versus damages; guilt versus sentencing).

Relatedly, there are also many more possible ways of addressing statistical evidence than merely rejection as inadmissible on one hand or unqualified acceptance on the other. Possible intermediate rules and mechanisms include restricting admissibility of statistical evidence to civil trials (or to the sentencing phase of the criminal trial); treating statistical evidence as admissible across civil and criminal trials but prohibiting its use as sole basis for a positive finding or for imposition of liability (whether civil or criminal); and treating statistical evidence as categorically admissible and as a sufficient basis for a finding of civil liability (or for sentencing purposes) but prohibiting its use as sole basis for conviction. These intermediate categories are not conclusive, of course, and do not exhaust the full range of possibilities for regulating the use of statistical evidence. Further distinctions can and will be drawn, such as those differentiating between exonerating and incriminating statistical evidence, or between different requirements for corroborative evidence, and so forth.

Our discussion in the rest of this section will be, to repeat, very preliminary. To exemplify the potential of applying the analysis in earlier parts of this paper to evidence law doctrines, we briefly discuss two central cases – DNA evidence and predictive evidence at the conviction stage.

DNA

DNA is an interesting case for illustrating the applicability of our theory to the legal arena. It is interesting because DNA evidence seems statistical, but the law seems rather happy to rely on it, its misgivings about statistical evidence notwithstanding. In the words of one court, DNA testing is the "single greatest advance in the search for truth . . . since the advent of cross-examination"⁷². And to most, this seems like a reasonable legal policy here. Can this be explained? Can what we say about statistical evidence in previous sections shed some light on this notable exception⁷³?

We need to make the following three preliminary points. First, we are of course not going to do the relevant science. We are going to assume that DNA is highly probative. Indeed, we are going to restrict our attention to just those cases where, say, the probability that the accused is guilty given that there is a DNA match, is extremely high, though not 1. Second, there are many ways of using DNA evidence. We are going to focus our attention on the hard case of so-called cold-hit DNA evidence, where DNA is the only evidence, and it was obtained without some prior suspicion – we will be assuming, in other words, that a DNA sample was obtained, run against some database, and a match was found, not that a suspect was pinpointed, then tested for DNA. And third, we will be restricting our attention to just the use of DNA evidence as evidence for the prosecution (in a criminal case). At least in the criminal case, the legitimacy of DNA evidence as exonerating evidence is clear enough not to be interesting (the relevant high probability of accuracy certainly suffices for reasonable doubt).

With these assumptions in place, then, can anything be said in favor of using DNA evidence, even against a background of suspicion towards statistical evidence in general? Well, one obvious special feature of DNA evidence compared to most other kinds of statistical evidence is the extremely high probabilities relevant here. So it is natural to think that this is what makes a difference here. But why, exactly?

⁷² **People v. Wesley**, 533 N.Y.S.2d 643, 644 (N.Y. Sup. Ct. 1988).

⁷³ It is not, of course, the only exception. But it is probably the clearest one.

A possible – reasonable, but unexciting – answer may be in terms of the relative value of accuracy here. It may be thought, that is, that the same reasons we always have not to rely on statistical evidence – whatever exactly they are – are fully present here. It's just that here, because of the extremely high probabilities, the significance of the value of accuracy is much weightier than in other cases, and so in this case the reasons related to the value of accuracy outweigh the standard reasons not to rely on statistical evidence.

Or, one may try to tell an incentive-story here. Recall (the generalization of) Sanchirico's theory, according to which relying on statistical evidence will create inefficient incentives for, say, the Blue Bus Company as well as its competitor the Red Bus Company. Sanchirico's reasoning relied on the fact that both companies will be in a position to know that their chances of being found liable are independent of their relevant conduct (because liability is determined by their market share). But perhaps in DNA cases – certainly, in *most* DNA cases – such knowledge is not available. So the incentive story arguably does not apply. So we do not have the incentive-based reason to ignore genuinely probative statistical evidence. This line of thought sounds reasonable enough, but we are not sure that it captures the full extent of our intuitions here. Suppose, for instance, that except for DNA, we can also check for a DNA* match. DNA* shares with DNA its incentives-relevant properties (things like what knowledge is and is not available *ex ante*), but is much less effective probabilistically, so that the probability that the accused is guilty given a DNA*-match is, say, around 70%. In such a case too the incentive story doesn't go through. But our intuitive reluctance to relying on statistical evidence is still fully present⁷⁴.

Another line of thought starts with the suggestion briefly mentioned in section 1 above, according to which at least one of the things we find problematic about statistical evidence is that systematic reliance on it guarantees some false decisions, indeed false convictions (think

⁷⁴ Perhaps, then, we should think of the incentive story as that which justifies relying on DNA evidence, and offering a debunking explanation of our intuitions about DNA*-evidence.

here of a variant of the gatecrasher case, where we indict all those attending the stadium). This problem – a guaranteed mistaken decision – doesn't seem to be relevant to DNA evidence, where systematically relying on such evidence doesn't have a similar guaranteed result. But this explanation too does not do all the needed work here. This is partly because of the doubts mentioned in section 1 (roughly, why think that the difference between a guarantee and a ridiculously high probability that is still smaller than 1 makes all this difference?), but also for the following reason: We can imagine a variant of the gatecrasher case where the guarantee of a false conviction is absent – say, if many of those attending the stadium have escaped before the police arrived. Indeed, suppose (again, as we did in section 1) that only one person from the stadium was apprehended, and only he will be tried. Then relying on statistical evidence does not have the result of a guaranteed false conviction. But the reluctance to rely on statistical evidence is still very much present. So the guarantee story can't be the full story here⁷⁵.

The stories just sketched hold, without a doubt, at least some of the relevant truth about DNA evidence. And we in no way want to deny that – in fact, we are committed to such non-epistemic stories being the ones that should guide legal policy. But we want to show – if only briefly, and somewhat speculatively – a possible relation here to the epistemic considerations highlighted in earlier sections. Suppose, then, that we convict A solely because of a DNA-match. Had A not been guilty, would we have still convicted him? Well, had A not been guilty, but had the DNA evidence still matched A's, we would have still convicted him. But this is a different counterfactual – in the terms of section 5, it invites us to travel to a different possible world. The counterfactual relevant here is the one we started with, and here there's considerable pressure to answer in the negative: Had it not been A, we (most probably) wouldn't

⁷⁵ In general, it is an interesting exercise to construct a parallel gatecrasher case for any story about DNA evidence. For instance, in the case of DNA we typically don't even know – in a specific case – that there is another person whose DNA would match that found in the crime scene. So perhaps we should think about a gatecrasher case where we don't know that some people actually bought tickets; all we have is the probability that some did. Etc. Things get complicated, and as they do it becomes less clear what exactly to say. But even in this last kind of gatecrasher case it seems that the law wouldn't convict just on the grounds of the statistical evidence, and it also seems (though to two of us more than to the third) that this is as it should be. So DNA remains special.

have found a DNA-sample matching A's in the crime scene. In possible-world talk – there's considerable intuitive pressure to think that a possible world in which A is innocent and yet the DNA sample matches A's is farther from the actual world than a world in which A is innocent and no matching DNA sample is obtained at the crime scene. If this is true, then in the DNA case – unlike other statistical evidence cases – Sensitivity is satisfied. So DNA may be special even epistemically, according to the account in this paper.

The same point can also be put in terms of the explanatory test introduced in section 5. If we convict someone of gatecrashing just on the strength of the statistical evidence, and later find out that she was a rare ticket-buyer, we do not (nor does it seem that we should) look for a deep explanation – we played the odds, and lost. But in a case where we convict A based purely on DNA evidence and later we find out he was innocent, we do look for a deeper explanation, and justifiably so, it seems: Such mistake does call for explanation⁷⁶.

Of course, much more needs to be said here. First, as you may recall, we did not have much to say about the proximity relation between worlds. For many cases, this problem remains theoretical. But now a theory of this proximity relation seems more badly needed, as the point in the paragraph before last seemed to rely on rather intricate characteristics of that relation. Unfortunately, though, we do not have such a theory up our sleeves. Let us just note, then, that our story here does depend on substantive assumptions about the proximity relation, and that at least some accounts of it work nicely with the points just made⁷⁷. Second, even if we acknowledge that there is this Sensitivity-relevant difference between DNA and other statistical evidence, we can still ask what it is about DNA that renders it special in this way. It seems natural to again rely on the high probability here. But the relations between probabilities,

⁷⁶ Perhaps we would say (in this case as apposed to the gatecrasher case) that the possibility of foul play - that someone tampered with the DNA sample - becomes relevant (other things being equal).

⁷⁷ Thus, David Lewis (1979) gives a contextual account of the proximity relation in terms of the number and size of miracles needed to move from the actual world to the relevant possible world. And this fits nicely with the points in the text – it would seem like a fairly big miracle for A to be innocent and yet for the DNA sample from the crime scene to match A's. See Bennett (2003) for some criticism of Lewis's account.

proximity relations between worlds, and indeed what does and what does not call for explanation are anything but simple or obvious⁷⁸.

Still, we hope to have succeeded in at least sketching some practical considerations that make DNA special, as well as some related epistemic considerations that go well side-by-side with them.

Predictive Evidence at the Conviction Phase of Trial

X is a 19 year old man, raised in a dysfunctional family, facing charges of assault and battery. Let's suppose that reliable statistical evidence points to the fact that young males who were raised in broken homes have a significantly greater than average propensity for violent crimes (say, 65% of such individuals have been convicted of violent crimes). Could this evidence be admitted as proof that X is implicated in the act with which he is currently charged on the particular occasion of his trial? ⁷⁹ Under prevailing doctrine, the answer to this question is "No". This type of evidence would either be deemed irrelevant or inadmissible at the conviction phase of trial. Our analysis can explain why.

The first thing to note here is that Sanchirico's incentive-based story straightforwardly applies. Because the nature of the predictive evidence is already set when the agent deliberates on whether or not to commit the offense (at that point in time, the 19-year-old cannot change his age, gender, or the fate of his parents' marriage), relying on such evidence will have counter-productive effect on the relevant incentives for primary behavior.

Now, there's a complication here: Sanchirico's analysis focuses on the legal payoff in the period of time following the act suggesting bad character, or – in the case of prior convictions – following the involvement in the first offense. But it is not clear why, when devising rules for

⁷⁸ Recall (note **) that in the lottery cases, typically (but not always), the probability that a given ticket will win is the same as any other ticket. Not so for different mistakes a newspaper might make when reporting a wrong lottery winner (and equally for the other individual evidence cases).

⁷⁹ For further discussion of this hypothetical see Redmayne's (2008, 281) "future violence" example.

optimal deterrence, focus should be placed exclusively on the incentive structure and on the legal payoff in that period of time. After all, individuals should also be deterred from committing the first act of crime. And when focusing on the legal payoff prior to the first crime, a rule that permits information of prior convictions to be admitted to the court may actually further deterrence. It would enhance the expected cost of engaging in the first criminal act, for the expected sanction would now include a greater probability of conviction in any future trial.

Interestingly, though this may be a problem for Sanchirico's analysis in some contexts, it does not apply in the kind of case discussed in this section – namely, evidence of propensity for crime based on gender, age, ethnic or demographic characteristics. Unlike the feature of engagement in criminal activity underlying the prior convictions category, characteristics such as age or gender cannot be chosen by the individual nor are they problematic from a social welfare perspective. They are not, in other words, to be deterred. So the focus on just the point in time in which the agent deliberates on whether or not to commit the offense for which he is now on trial – though perhaps problematic in some contexts – is justified in ours.

Here too, the incentive-based analysis can be complemented by the kind of epistemological story we have been telling. The first thing to note here is that the predictive evidence of the kind mentioned seems – even if very strong – not to support knowledge. We seem to be strongly disinclined to claim or attribute knowledge (that so-and-so will commit a violent offense, or even that he has) on the basis of such evidence alone. We are even hesitant to claim justified belief on such ground. Sensitivity-style counterfactuals nicely show why, of course, because such predictive evidence is paradigmatically insensitive.

The explanatory test, though, is a more interesting and complicated case. Against the background of the grim statistics – a very high percentage of young males from broken homes are guilty of violent crimes, say – does the fact that a specific young male with such background is not involved in such violence call for explanation? Or do we tend to take a you-win-some-you-lose-some attitude towards it? The answer, it seems to us, is not clear. If someone succeeds in

living a well-adjusted, law-abiding life coming from a background where most people do not, this does seem to call for explanation – perhaps the kind of explanation that shows that he’s not after all a typical representative of that background (he had a loving and nurturing grandparent, an especially resilient character, etc.). But it’s not clear what to make of this: First, this may show that we tend to reject the general predictive evidence here not because it’s statistical, but rather because it’s not good statistical evidence (as the youngster we’re interested in is not representative of the relevant reference class). Second, our insistence on finding this kind of explanation here may be partly attributed to wishful thinking, and for this reason debunkable. Third, if we do go ahead and make a decision that partly relies on the predictive evidence, and we then find out that it misled us, we will be inclined to take a you-win-some-you-lose-some attitude – if not towards the youngster’s success directly, at least towards our procedure of making the relevant decision. So while we acknowledge that the explanatory story here is not as clear-cut as the knowledge one or as the one in terms of Sensitivity, still it doesn’t strongly pull in the opposite direction either.

All of this has been, to repeat, very preliminary, and much more needs to be said here if we are to show that our vindication of the general suspicion towards statistical evidence – the dual story, first in epistemological terms, then in incentive-based terms – can successfully explain the many contours of evidence law doctrines (or can plausibly recommend reform when appropriate). But we hope the brief discussion in this section still gives some grounds for optimism about this project, which we hope to engage in more detail elsewhere. And if it does, then this optimism itself further supports the plausibility of the theoretical story in the earlier parts of this paper.

References

- Adams, E. W. (1970). Subjunctive and Indicative Conditionals. *Foundations of Language* 6: 89-94.
- Bennett, Jonathan (2003). *A Philosophical Guide to Conditionals*. Oxford: Clarendon Press.
- McBride, Mark (2011). Reply to Pardo: Unsafe Legal Knowledge. *Legal Theory* 17: 67-73.
- L.S. Carrier (1971). An Analysis of Empirical Knowledge. *Southern Journal of Philosophy* Vol. 9: 3-11.
- Colyvan, Mark ; Regan, Helen M. & Ferson, Scott (2001). Is it a crime to belong to a reference class. *Journal of Political Philosophy* 9 (2):168–181.
- Conee, Earl & Feldman, Richard (1998). The Generality Problem for Reliabilism. *Philosophical Studies* 89: 1-29.
- DeRose, Keith (1995). Solving the Sceptical Puzzle. *Philosophical Review* 104: 1-52.
- DeRose, Keith (2010). Insensitivity is back, baby! *Philosophical Perspectives* 24 (1):161-187.
- Dretske, Fred (1970). Epistemic Operators. *The Journal of Philosophy* Vol. 67: 1007-1023.
- Dretske, Fred (1971). Conclusive Reasons. *Australasian Journal of Philosophy* 49: 1-22.
- Fischer, David A., "Successive Causes and the Enigma of Duplicated Harm", 66 TENN. L. REV. (1999) 1127.
- Goldman, Alvin (1976). Discrimination and Perceptual Knowledge. *The Journal of Philosophy* Vol. 78: 771-791.
- Goodman, Nelson (1947). The problem of Counterfactual Conditionals. *The Journal of Philosophy*, Vol. 44.
- Hawthorne, John (2004). *Knowledge and Lotteries*. Oxford: OUP.
- Hawthorne, John & Lasonen-Aarnio, Maria (2009). Knowledge and Objective Chance. in *Williamson on Knowledge*. P. Greenough & D. Pritchard (eds.). Oxford: OUP.
- Hájek, Alan (MS). Most Counterfactuals Are False. Monograph in progress:
http://philosophy.cass.anu.edu.au/sites/default/files/Most%20counterfactuals%20are%20false.1.11.11_0.pdf
- Haw, Rebecca (2009). Prediction Markets and Law: A Skeptical Account. 122 Harv. L. Rev. 1217, 1229
- Hock Lai Ho, (2008) *A Philosophy of Evidence Law* (2008), Oxford: Oxford University Press.
- Kaye, David, "The Paradox of the Gatecrasher and other Stories" ARIZONA ST. L. J. 101 (1979)
- Koehler, Jonathn J. & Shaviro, Daniel, "Veridical Verdicts: Increasing Verdict Accuracy through the Use of Overtly Probabilistic Evidence and Methods", 75 Cornell L. Rev. (1990) 247
- Lempert (2001) ...
- Lillquist, Eric "Recasting Reasonable Doubt: Decision Theory and the Virtues of Variability", 36 U.C. Davis 85 (2002)
- Lewis, David (1973). *Counterfactuals*. Oxford: Basil Blackwell.
- Lewis, David (1979). Counterfactual dependence and time's arrow. *Noûs* 13 (4):455-476.

- Lewis, David (1986). *On the plurality of worlds*. Oxford: Blackwell.
- New, Christopher (1992). Time and Punishment. *Analysis* 52: 35-40
- Nozick, Robert (1981). *Philosophical Explanations*. Cambridge: Harvard University Press.
- Pardo, Michael S. (2010). The Gettier Problem and Legal Proof. *Legal Theory* Vol. 16.
- Pardo, Michael S. (2011). More on the Gettier Problem and Legal Proof. *Legal Theory* Vol. 17: 75-80.
- Pritchard, D. H. (2005). *Epistemic Luck*. Oxford: Oxford University Press.
- Pundik (2008) ...
- Pundik's (The Inevitable Efficiency of Using Racist Statistical Evidence in Court)...
- Mike Redmayne (2008). Exploring the Proof Paradoxes. *Legal Theory* 14: 281-309.
- Rhee, Robert J "Probability, Policy, and the Problem of Reference Class", INT'L J. OF EVIDENCE & PROOF, 11 Evidence & Proof (2007) 286.
- Alex Stein, (2005) *Foundations of Evidence Law* Oxford: Oxford University Press
Chris Sanchirico (2001). Character Evidence and the object of Trial. 101 COLUM. L. REV. 1227.
- Schoeman, Ferdinand (1987). Statistical vs. Direct Evidence. *Noûs* 21: 179-198.
- Sharon, Assaf & Spectre, Levi (forthcoming). Evidence and the Openness of Knowledge. *Philosophical Studies*.
- Sharon, Assaf & Spectre, Levi (MS). Three Problems for Williamsonian Epistemology.
- Sider, Theodore (2010). *Logic for Philosophy*. Oxford: Oxford University Press.
- Smilansky, Saul (1994). The Time to Punish. *Analysis* 54 (1):50 - 53.
- Smith, Martin (2010). What Else justification Could Be. *Noûs* 44: 10-31.
- Sorensen, Roy (2006). Future Law: Prepunishment and causal theory of verdicts. *Noûs* 40: 166-183.
- Stalnaker, Robert (1968). A Theory of Conditionals. in *Studies in Logical Theory, American Philosophical Quarterly*. Oxford: Blackwell. Monograph Series 2: 98-112.
- Statman, Daniel (1997). The time to punish and the problem of moral luck. *Journal of Applied Philosophy* 14 (2):129–136.
- Thomson, Judy (1986). Liability and Individualized Evidence. in *Rights, Restitution, and Risk* William Parent (ed.): 225–250. Cambridge: Harvard University Press.
- Vogel, Jonathan (1990). Are There Counterexamples to the Closure Principle? in *Doubting: Contemporary Perspectives on Skepticism*. M. Roth and G. Ross (eds.). Dordrecht: Kluwer.
- Wasserman, "The Morality of Statistical Proof and the Risk of Mistaken Liability" 13 *Cardozo Law Review* 935 (1991).

Williamson, Timothy (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.

Williamson, Timothy (2009). Reply to John Hawthorne and Maria Lasonen-Aarnio. in *Williamson on Knowledge*. P. Greenough and D. Pritchard (eds.), Oxford: OUP.

Williamson, Timothy (Forthcoming). Improbable Knowing. in T. Dougherty (ed.) *Evidentialism and its Discontents*. Oxford: Oxford University Press.

Wright, Richard, "Causation, Responsibility, Risk, Probability, Naked Statistics, and Proof: Pruning the Bramble Bush by Clarifying the Concepts", 73 *Iowa Law Review* (1988) 1000, 1050.