

DAVID ENOCH,
LEVI SPECTRE, AND
TALIA FISHER

Statistical Evidence, Sensitivity, and the Legal Value of Knowledge

I. THE PROBLEM

A bus causes harm. In the first scenario, an eyewitness recognizes the bus as belonging to the Blue Bus Company. The witness, however, is imperfectly reliable; let us say that she is roughly 70 percent reliable in matters such as this one. The law has no qualms about accepting the eyewitness 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 eyewitness, but we have uncontested data regarding the distribution of buses in the relevant area; in particular, the Blue Bus Company owns roughly 70 percent 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

Previous versions of this article were 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; at the Theoretical and Practical Rationality group at the Institute for Advanced Studies in Jerusalem in the fall of 2011; at the Inter-Disciplinary Center, Herzeliya, Faculty of Law Colloquium in the fall of 2011; at the New England Consequentialism Workshop in February 2012; and at the Economic Theory Workshop at Tel Aviv University in March 2012. We thank the participants for the valuable discussion. We especially thank our commentator on that first occasion, Amit Pundik, for very detailed and helpful comments. Thanks to Hagit Benbaji, Selim Berker, Mitch Berman, Stewart Cohen, Cian Dorr, Kate Elgin, Nir Eyal, Yuval Eylon, Johann Frick, Caspar Hare, Amnon Knoll, Netanel Lipshitz, Ofra Magidor, Ofer Malcai, Jennifer Nagel, Christian Piller, Ariel Porat, Jim Pryor, Mike Redmayne, Melinda Roberts, Sergio Tenenbaum, Ruti Weintraub, Dan Wikler, and two referees and an editor for *Philosophy & Public Affairs* for helpful comments and discussions. David Enoch's research was supported by the Israel Science Foundation (grant no. 136/09).

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 eyewitness 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 equivalent in both cases. And after all, there is *something* statistical about the eyewitness testimony as well.³ She is fallible, of course, and out of an arbitrary set of, say, a hundred buses that she recognizes as Blue Bus buses in similar circumstances, about thirty will belong to other companies.⁴ In what way, then, can we distinguish between the fact that the relevant bus was a member of a set of which 70 percent belonged to the Blue Bus Company (not enough to ground a finding), and the fact that the bus was a member of a set of which 70 percent are accurately identified by the eyewitness (sufficient to ground a finding)?⁵ How can we, then, accommodate in a theoretically respectable way the evidential distinction between the two scenarios we started with?⁶

1. See, e.g., Rebecca Haw, "Prediction Markets and Law: A Skeptical Account," *Harvard Law Review* 122 (2009): 1217–38, at pp. 1217, 1229, and the references there.

2. For one discussion of this common hypothetical, see Mike Redmayne, "Exploring the Proof Paradoxes," *Legal Theory* 14 (2008): 281–309.

3. See, e.g., Jonathan J. Koehler and Daniel Shaviro, "Veridical Verdicts: Increasing Verdict Accuracy through the Use of Overtly Probabilistic Evidence and Methods," *Cornell Law Review* 75 (1990): 247–79, at p. 263.

4. Nothing in our relevant intuitions depends on the nonstatistical evidence being witness testimony. Think about other types of clearly individual evidence, like the videotape of a surveillance device, or some physical evidence (say, that the accused's gun was found on the scene).

5. Sometimes people write as if individual evidence is *about* the individual in a way that statistical evidence is not. (See, e.g., Richard Wright, "Causation, Responsibility, Risk, Probability, Naked Statistics, and Proof: Pruning the Bramble Bush by Clarifying the Concepts," *Iowa Law Review* 73 [1988]: 1001–77, at p. 1050.) We have no objection to this way of putting the initial intuition here, but it is important not to mistake a new name for the problem for a solution. For it remains to be explained in what way exactly individual evidence is—and statistical evidence is not—about the individual. Surely, both are *indicative* of something about the individual. Is this not enough? (See Ferdinand Schoeman, "Statistical vs. Direct Evidence," *Noûs* 21 [1987]: 179–98, at pp. 183–84, for a somewhat similar point.)

6. One major issue we will not try to tackle here is the reference class problem. We will assume, more or less throughout, that the statistical evidence latches on to the relevant

The problem generalizes, of course. We can easily think of examples in criminal law or in the law of torts (or elsewhere, such as in other areas of the law or perhaps in determining public policy and so on), examples with different levels of probability (that are held constant in both scenarios), examples where there is or there is not 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 examples.⁷ 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 is to declare the law of evidence—and indeed, the powerful, uncompromising intuitions of pretty much all of its scholars⁸—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 possibility,⁹ it seems like it should be avoided if possible. We are, then, going to make an attempt to find after all a vindicating explanation of this distinction. Note that we will not here attempt a vindication of *some specific way* of discriminating between statistical and individual evidence (say, in terms of admissibility). The initial challenge is to vindicate *any distinction at all*, to show any reason at all to treat statistical evidence more suspiciously than individual evidence (as the law seems to do, across jurisdictions).

In the next section, we distinguish between an epistemic and a practical way of addressing the initial challenge. 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

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 percent reliability of the eyewitness).

7. For a list of examples, and some insistence on the significance of the differences between them, see Redmayne, "Exploring the Proof Paradoxes," pp. 282–85.

8. See Alex Stein, *Foundations of Evidence Law* (Oxford: Oxford University Press, 2005).

9. Schoeman, "Statistical vs. Direct Evidence," 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.

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. But we then argue (in section 5) that such epistemological niceties as whether a finding is sensitive are not something the law should care about. Instead, we offer an instrumental, incentive-based argument for why the law should be suspicious of statistical evidence. This does not mean, though, that the epistemological story has been left behind. Rather, we show that there is a close connection between the incentive-based story and the epistemological story—both rely on Sensitivity-style counterfactuals. That the epistemic and instrumental considerations nicely cohere, then, is not merely a happy coincidence. A brief conclusion follows.

II. TWO KINDS OF SOLUTIONS

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

Instances of the first, the 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 is harder for statistical evidence than it is for (probabilistically equivalent) 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 individual evidence can support knowledge, but statistical evidence cannot. All of these are epistemological 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, avoiding the sandwiches from the deli with the questionable past track record. But justifying actions or rendering them rational is one thing, and justifying beliefs or rendering

them rational is arguably another. The epistemic strategy insists that there is a difference between statistical and individual evidence that is of this latter, epistemic kind.

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 practical kind.

Instances of the second, the 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 scenarios, the game—this game—is over. But the probability of truth 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.

With the distinction between epistemic and practical strategies¹⁰ (for vindicating the distinction between statistical and individual evidence) at hand, we can now make the following prediction: if the best vindication of this distinction is along practical—certainly, 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's so treating it, then it is essentially tied to the *law's* so treating it. It is, after all, quite possible 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

10. See *ibid.*, p. 187, for a similar distinction. Redmayne, "Exploring the Proof Paradoxes," p. 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 need not be limited to instrumental considerations.

the distinction, then it is likely to apply much more widely, indeed perhaps as widely as the relevant epistemic notion (justification, perhaps, or knowledge) reaches. In the other direction now, if the problem arises much more widely 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.

III. SENSITIVITY

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

Think about the following version of the lottery paradox (for knowledge):¹¹ In the first scenario, you buy a lottery ticket, where the chance of winning is (literally) one in a million. The winning ticket has been picked, but you receive no indication about the results. Do you know that your ticket has not won? The answer that seems overwhelmingly plausible is “no.” You may know that it is highly unlikely that your ticket won; you may be justified in gambling against your ticket at rather high odds; you may, of course, *believe* that it has not won; but you do not *know* that it has not won (and this even if, as things turn out, your ticket has not in fact won).

Compare this, now, to the following, second scenario. In this scenario you buy a lottery ticket with somewhat better odds—one in a thousand, perhaps. You hold on to it for a day. Now the winning ticket has 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 us 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 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.)

11. See John Hawthorne, *Knowledge and Lotteries* (Oxford: Oxford University Press, 2004); and H. E. Kyburg, *Probability and the Logic of Rational Belief* (Middletown, Conn.: Wesleyan University Press, 1961).

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 has not 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 constant¹²—what can possibly explain this difference?

One plausible answer is in terms of some relevant counterfactuals.¹³ 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 are 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, we are 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

12. This further condition may be needed in order to alleviate some internalist worries (in the sense “internalism” has in epistemological contexts); we do not think we need to worry about them here.

Some epistemologists (following Timothy Williamson, *Knowledge and Its Limits* [Oxford: Oxford University Press, 2000]; and Williamson, “Reply to John Hawthorne and Maria Lasonen-Aarnio,” in *Williamson on Knowledge*, ed. P. Greenough and D. Pritchard [Oxford: Oxford University Press, 2009], pp. 313–30) may want to make a distinction here between epistemic probabilities and chances. Others (following Stewart Cohen, “How to Be a Fallibilist,” *Philosophical Perspectives* 2 [1988]: 91–123) 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, as in Hawthorne, *Knowledge and Lotteries*). For our purposes here, though, we do not need to decide these issues.

13. Prominent proponents of this kind of view include Robert Nozick, *Philosophical Explanations* (Cambridge, Mass.: Harvard University Press, 1981); Keith DeRose, “Solving the Sceptical Puzzle,” *Philosophical Review* 104 (1995): 1–52; Alvin Goldman, “Discrimination and Perceptual Knowledge,” *Journal of Philosophy* 78 (1976): 771–91; and Fred Dretske, “Conclusive Reasons,” *Australasian Journal of Philosophy* 49 (1971): 1–22.

highly plausible to say that in the newspaper scenario, had the belief (that your ticket had not won) been false, you would not have held it. 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 proposition 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 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 epistemic fluke—you, after all, would have believed it even had it been false. And so it may be thought that there is 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 would not have believed it any longer), your belief does seem appropriately sensitive to the truth, you do seem entitled to some intellectual credit here, and 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: *S*'s belief that *p* is sensitive = *df.* Had it not been the case that *p*, *S* would (most probably)¹⁴ not have believed that *p*.

14. This qualification allows us to bypass some technical issues here, having to do with the common way of fleshing out counterfactual semantics in terms of possible worlds. (See, e.g., Robert Stalnaker, “A Theory of Conditionals,” in *Studies in Logical Theory*, *American Philosophical Quarterly*, Monograph Series 2 [Oxford: Blackwell, 1968], pp. 98–112; and David Lewis, *Counterfactuals* [Oxford: Basil Blackwell, 1973]; and for more recent discussions, see Theodore Sider, *Logic for Philosophy* [Oxford: Oxford University Press, 2010]; and Jonathan Bennett, *A Philosophical Guide to Conditionals* [New York: Oxford University Press, 2003].) On such proposals, roughly speaking, a counterfactual is true if and only if in the closest possible world in which its antecedent is true, so is its consequent. The problem is that worlds that are less likely to be the actual one (for instance, one where the newspaper made a mistake) are not guaranteed to be further from the actual one compared to more likely worlds. The most-probably qualification bypasses this problem.

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

Epistemologists have been discussing Sensitivity for several decades now, and many details have emerged. On the optimistic side, Sensitivity and related requirements have been shown to help with a fairly wide class of epistemological problems (including at least some Gettier cases, Barn Façade cases first discussed by Goldman,¹⁵ and perhaps also some skeptical puzzles). On the downside, though, clever counterexamples to Sensitivity have been devised and revisions offered, followed by other counterexamples, cleverer still, being devised and further revisions being offered, with the debate going on as lively today as it was thirty years ago.¹⁶ While our employment of Sensitivity is of course vulnerable to any general objection to Sensitivity, we think that we can avoid many of these problems here, relying on a division of labor among philosophers in the following way: We are going to rely only on Sensitivity as presented above, and we are only going to be assuming that something in its vicinity is of genuine epistemological importance.¹⁷ We are not, let us emphasize, going to insist that Sensitivity is a necessary condition for knowledge. We are going to insist merely on its (or something sufficiently close to its) being an epistemological desideratum: with other things held equal, a sensitive belief is epistemically better than an insensitive one. This will suffice for our purposes here, as we will see below. And we can thus let more purely epistemological texts develop the details here.

IV. SENSITIVITY AND STATISTICAL EVIDENCE

Surely, by now you will have noticed the similarity. The two Blue Bus scenarios (market-share evidence and eyewitness testimony) parallel the two lottery scenarios (where the belief is based merely on the odds and

15. Alvin Goldman, "Discrimination and Perceptual Knowledge," *Journal of Philosophy* 78 (1976): 771–91.

16. For two recent contributions, and for references, see Keith DeRose, "Insensitivity Is Back, Baby!" *Philosophical Perspectives* 24 (2010): 161–87; Saul A. Kripke, "Nozick on Knowledge," in *Philosophical Troubles: Collected Papers*, vol. 1, ed. Saul A. Kripke (Oxford: Oxford University Press, 2011), 162–224.

17. For one example of something in the vicinity of Sensitivity, see the discussion of normic support below.

where it is also based on the newspaper report, respectively).¹⁸ The parallel is not just based on the intuitive similarity (the cases do have a similar feel to them, do they not?). 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 us 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 eyewitness testimony—it seems like our finding is sensitive. Had it not been a Blue Bus bus, would we have found the Blue Bus Company liable? Probably not. Our eyewitness 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 percent

18. The evidence-law literature on statistical evidence (and 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, *Foundations of Evidence Law*, p. 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, “Exploring the Proof Paradoxes,” p. 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) lottery paradox, but he fails to mention the relevance of something like Sensitivity (rather, he discusses the Safety condition, which is related but less appropriate here, and even that he discusses only in a very sketchy way). Hock Lai Ho, *A Philosophy of Evidence Law* (Oxford: Oxford University Press, 2008), pp. 168–69, briefly mentions the similarity but fails to put it to theoretical use—see here also note 27. Judy Thomson, “Liability and Individualized Evidence,” in *Rights, Restitution, and Risk*, ed. William Parent (Cambridge, Mass.: Harvard University Press, 1986), pp. 225–50, at pp. 236–39, does notice the parallel and makes more of it; indeed, there are some important similarities between her treatment of statistical evidence and ours. We hope to compare her view to ours in more detail elsewhere. Let us just note here that Thomson’s discussion is located in the context of a causal theory of knowledge, which we (and pretty much everyone else these days) reject, and that she seems to assume that the law should significantly care about knowledge and epistemology, a point we emphatically deny in this article.

of the buses; it is just that the bus that causes the harm is no longer one of its buses (rather, it is a Red Bus bus). In such a case, we would still have the exact same statistical evidence available to us. So in that case, too, we would have found the Blue Bus Company liable. So by relying on statistical evidence, we render our findings insensitive.

Or consider the gatecrashers case,¹⁹ 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, even when convicting on the strength of a probabilistically equivalent piece of direct individual evidence (say, a videotape) seems perfectly fine. And now we can explain, it seems, why: Even if the percentage of gatecrashers among all those attending the stadium is quite high, 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 still have 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 practical 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 in any way 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

19. David Kaye, "The Paradox of the Gatecrasher and Other Stories," *Arizona State Law Journal* (1978): 101–43, at p. 104; Robert J. Rhee, "Probability, Policy, and the Problem of a Reference Class," *International Journal of Evidence and Proof* 11 (2007): 286–91, at p. 289.

very same reluctance.²⁰ 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 the epistemic significance of Sensitivity explains why there is something suspicious about statistical evidence across the board—in the legal context and in the lottery cases, as well as anywhere else where we care about knowledge or about our beliefs' being sensitive to the facts.

Another, closely related way of offering an epistemological vindication of the kind needed here is explanatory. 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 percent 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 percent of the time, and that is that: tough luck this time. But when individual evidence misleads, the situation is different.²¹ If we relied on the eyewitness in ruling against the Blue Bus Company, and it turns out that the bus that caused the harm actually belonged to the Red Bus Company, this discrepancy seems to call for explanation. Certainly,

20. Amik Pundik, "Statistical Evidence and Individual Litigants: A Reconsideration of Wasserman's Argument from Autonomy," *International Journal of Evidence and Proof* 12 (2008): 303–24, at p. 303, emphasizes that in some contexts—say, that of supplying a medical diagnosis—we are perfectly willing to rely on statistical evidence. He takes this as reason to believe that the problem is not as general as we say in this article that it is. But note that Pundik is interested in the extent to which we are willing to *act* on statistical evidence in different practical contexts. The point in this article is that even outside the context of any action at all, we are unhappy to attribute knowledge (and perhaps even justified belief) to a subject who bases her beliefs on statistical evidence alone. And here the reluctance seems to be very general indeed. We briefly revisit this point below.

21. 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 (for example, newspapers printing erroneous lottery results are not equally likely to print all erroneous results). Notice that such details—it seems reasonable to expect—will be closely connected to the requirement to explain the relevant kind of mistake that is emphasized in this article.

settling for a “you win some, you lose some” attitude seems out of place. In the terms Smith 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 for.²²

Now, given some plausible hypotheses about the relations between the truth-values of counterfactuals and the nature of the relevant kind of explanation, this explanatory story—in terms of normic support—seems closely related to our Sensitivity story. And in a fuller epistemological discussion we would have to ask which one (if any) is more basic, which one does the ultimate epistemological work here. But 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. 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 normic support, so long as something like Sensitivity is still epistemically relevant, our vindication of the distinction between statistical and individual evidence goes through.²³ At this point, this is partly a promissory

22. Martin Smith, “What Else Justification Could Be,” *Noûs* 44 (2010): 10–31, goes much further than merely noting the explanatory point in the text. He thinks of normic support as grounding epistemic justification, indeed as doing so even against probabilities, so that one belief may be justified while the other is not, even if the latter is more probable (for the thinker), so long as the former is 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 (“Reply” and “Improbable Knowing,” in *Evidentialism and Its Discontents*, ed. Trent Dougherty [Oxford: Oxford University Press, 2011], pp. 147–64) defending a sharp separation between chance and epistemic probability is relevant here. We do not need to discuss these interesting suggestions for our purposes here. At least one of us intends to elaborate—in future work—on some probabilistic differences that seem to mesh nicely with central cases of sensitivity and explanation.

23. Nevertheless, let us make the following points about Smith’s interesting suggestion (in “What Else Justification Could Be”). First, he explicitly addresses Sensitivity (p. 23), 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 eyewitness 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 eyewitnesses.) 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 that he is on to something important here, at

note to be made good on by the discussion in the next sections: we hereby officially promise that nothing in what follows will depend on Sensitivity's 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?) 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.

V. SHOULD THE LAW CARE ABOUT SENSITIVITY (OR KNOWLEDGE)?

Let us recap. Using (one version of) the lottery paradox, we introduced and motivated the intuitive requirement that beliefs should 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 probably 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 and presented 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.

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, even if Sensitivity is indeed necessary for knowledge, why should the law of evidence “care” about knowledge? Why, in other words, should it

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 Sensitivity-like counterfactuals—or the law-like connections that support them—are better candidates for being the more basic explanatory story.

make a legal difference whether a certain belief constitutes knowledge? In this section, we first present the remaining worry in more detail (in section 5A). We then present (in 5B) Sanchirico's incentive-based discussion of character evidence. Our solution to the puzzle—presented in section 5C—is going to concede that the law should not care about knowledge, or indeed about epistemology in general. And we are going to endorse a generalization of Sanchirico's instrumental reasoning as vindicating the difference in the law's attitude to statistical and individual evidence. But we will show that the relation between his instrumental reasoning and the epistemological discussion from the sections above is not a mere coincidence. The law should not care about epistemology, but it should care about something that is epistemically relevant: it should care about Sensitivity. Or so we are about to argue.

A. The Remaining Puzzle: Why Care about Knowledge?

It is important 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 mistakes and mistakes that are "too big." 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 parties') interests sufficiently improbable (of course, other considerations too may be relevant to determining the right procedures).²⁴

24. It is sometimes suggested that relying on statistical evidence amounts to preferring better consequences over doing justice in the specific case (e.g., Erik Lillquist, "Recasting Reasonable Doubt: Decision Theory and the Virtues of Variability," *U.C. Davis* 36 [2002]: 85–198, at p. 140), or to punishing people for being members of a reference class (e.g., Mark Colyvan, Helen M. Regan, and Scott Ferson, "Is It a Crime to Belong to a Reference Class?" *Journal of Political Philosophy* 9 [2001]: 168–81). Both of these suggestions are wrong, though, for reasons related to the point in the text: Individuals seem entitled not to be subjected to procedures that are too inaccurate, and statistical evidence improves accuracy. And even a court that is only concerned with making justice in the specific case has to use some evidence—the question is whether, when the court relies on evidence, statistical evidence is not as good as other evidence. And of course, whatever the evidence is that the court relies on for conviction, it punishes the convicted for the crime, not for being subject to incriminating evidence about the crime.

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 article, 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 or about epistemology more generally? 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?

Let us emphasize that to insist that the law should after all care about knowledge is (pretty much) to be willing to pay a price in accuracy. Indeed, excluding statistical evidence amounts to excluding (what is often) good, genuinely probative evidence. And this means that 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.

Perhaps the following thought experiment can serve to drive this point home: Suppose you have to choose the (criminal) legal system under which your children will live, and you can choose only between systems A and B. System A is epistemologically better: perhaps its courts only convict when they know (or think that they know) the accused is guilty, or perhaps they only convict based on sensitive evidence, or perhaps they convict only based on evidence that normically supports the conclusion that the accused is guilty. System B is not as good epistemically as System A. But System B is more accurate, so that the chances of System B convicting an innocent are lower than the chances of System A doing so.²⁵ Which system do you choose for your children:

25. We can put things here also in terms of luck, but it is not clear this will promote understanding rather than confusion. In a case of the kind described, in which system are *right* decisions right in a way that does not (or that does to a lesser extent) depend on luck?

the Epistemologically-Fine-But-Not-That-Reliable System A, or the More-Reliable-But-Not-That-Epistemically-Respectable System B? It seems to us that choosing System A amounts to an objectionable kind of epistemological fetishism, which disregards how much more accurate System B is. Now, this does not establish that knowledge (and the like) have no legal value whatsoever. It is consistent with this judgment that knowledge does perhaps have legal value, but a value that is always lexically inferior to that of accuracy. Though we suspect that knowledge has no legal value at all, for our purposes here we are happy to settle for the somewhat weaker claim: to the extent that it has value, its value is lexically inferior to that of accuracy.

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 is rational to have some belief given some evidence, or whether we are entitled to infer certain propositions from certain others. 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 attention.²⁶ 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 mistakes.²⁷

There is, we concede, a sense in which there is something about System A that makes it less vulnerable to luck (when it is mistaken, for instance, its mistakes call for an explanation and do not invite merely a “you win some, you lose some” kind of response). But when it comes to the law, it seems just as ill advised to prefer this over accuracy as it is to prefer knowledge or any other positive epistemic status.

26. See, e.g., Duncan H. Pritchard and John Turri, “The Value of Knowledge,” in *Stanford Encyclopedia of Philosophy*, <<http://plato.stanford.edu/entries/knowledge-value/>>, and the references there.

27. Notice that this remains so even if we engage in “knowledge-first” epistemology (see Williamson, *Knowledge and Its Limits*), perhaps partly because of a (purported) constitutive relation between assertability and knowledge. This is why Ho’s way (in *Philosophy of Evidence Law*, pp. 140–43) of addressing the proof paradoxes in the law of evidence seems

The point applies equally to the explanatory suggestion above: 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 storytelling otherwise, at least—just as harmful to the relevant party, just as detrimental to the relevant social interests, and so on, as mistakes that do call for explanation.

Our point is not that the law—not even evidence law—should “care” only about accuracy. Sometimes, other considerations can presumably trump accuracy. This is true in general,²⁸ and it may very well be true in our context as well. Perhaps, in other words, there are some cases of statistical evidence where other considerations trump accuracy: perhaps some profiling cases are of this kind. Our point is merely that *epistemological* considerations never by themselves seem to justifiably defeat considerations of accuracy when it comes to legal policy. At the very least, the legal relevance of epistemological considerations seems questionable enough to justify an attempt to see whether the distinction between statistical and individual evidence can be vindicated without relying on them.

to us unsatisfying. Even given his Williamsonian assumptions, why should the law care about, say, assertability? 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 opposed 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 John Hawthorne and Maria Lasonen-Aarnio, “Knowledge and Objective Chance,” in Greenough and Pritchard, *Williamson on Knowledge*, pp. 92–108; Williamson, “Reply”; and Assaf Sharon and Levi Spectre, “Epistemic Closure under Deductive Inference: What Is It and Can We Afford It?” *Synthese* (forthcoming).

28. See here, e.g., Mitch Berman's interesting discussion of how the fact that the standard used to allow instant replay to reverse the initial call in many sports requires indisputable or conclusive evidence is best explained not by a concern for accuracy, but for other values. Mitchell N. Berman, “Replay,” *California Law Review* 99 (2011): 1683–744.

There is another possibility to be considered here.²⁹ Epistemological considerations, we have just insisted, never *by themselves* defeat considerations of accuracy. But for anything thus far said, it is possible that epistemological considerations defeat considerations of accuracy *indirectly*, via some other considerations to which they are relevant. Perhaps, for instance, some *moral* judgments or attitudes—like blame—must be based on knowledge and so have to be sensitive in the sense above; and perhaps there are some close normative connections between legal findings and these moral judgments and attitudes (so that, say, we ought not to make a certain kind of a legal finding unless the relevant kind of moral attitude is appropriate). If so, epistemology may yet be allowed some significant role in the law. But too many details have to be filled in here (What moral judgments and attitudes? What kinds of legal findings? Why think that the appropriateness of the relevant moral attitudes depends on epistemological considerations? And why think that it is necessary for the appropriateness of the legal finding?), and it is not clear to us whether they can be filled in in a plausible way. We hope to address some of these issues in future work. For now, we will have to leave this work for the defenders of the (indirect) legal importance of epistemology. Until such work is done, we can rather safely (if somewhat tentatively) stick to the claim that epistemological considerations never defeat considerations of accuracy when it comes to desirable law.

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, as 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 a basis for legal policy.

29. Perhaps—we are not sure—such considerations can shed some light on the recent discussion of Gettier cases in the context of legal proof. See Michael S. Pardo, “The Gettier Problem and Legal Proof,” *Legal Theory* 16 (2010): 37–57; Pardo, “More on the Gettier Problem and Legal Proof,” *Legal Theory* 17 (2011): 75–80, and the references there.

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.³⁰ To see this, we need to discuss Sanchirico's work on character evidence.

B. Sanchirico on Character Evidence

Criminal law has mixed feelings about character evidence. Such evidence is typically admitted in the sentencing stage, but usually (at least traditionally) not as evidence for conviction; and this is the case despite the underlying suspicion that there too it can be, in a sense, good evidence and that it can serve to make courts' decisions more accurate. Furthermore, this mixed attitude toward 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 article, Chris Sanchirico concedes (for the sake of argument, at least) that there is no way of vindicating this attitude toward character evidence if we think of the law of evidence as exclusively aimed at helping courts in finding the truth or making factually accurate decisions.³¹ But Sanchirico suggests that we change the way we think of evidence law here. We should think of evidence law as being also—perhaps primarily—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 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 are already given. The character evidence that will be entered as evidence against (or for) him does not depend on his decision about 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

30. This is a different kind of indirect relevance from the one highlighted in the previous paragraph. There, it was indirect relevance by transitivity. In our case, it will be indirect relevance by a kind of parallelism.

31. Chris Sanchirico, "Character Evidence and the Object of Trial," *Columbia Law Review* 101 (2001): 1227–1311.

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—such as which is more relevant for deterrence—perhaps this line of thought can begin to vindicate the above-mentioned mixed attitude toward character evidence.

Of course, a lot may be going on with character evidence. And it need not 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, Sanchirico has still succeeded in drawing attention to another kind of consideration, one that it would be foolish for a legal system to ignore.

Sanchirico's article is not about statistical evidence, but about character evidence. There may be some similarities between the two, but we will not pursue them here.³² The important point for our purposes is that Sanchirico's general strategy can be easily applied to statistical evidence as well. Think, for instance, about John, a potential gatecrasher who is now deliberating, considering either purchasing a ticket, or perhaps gatecrashing, or perhaps 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 event at 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 percent of those attending are

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

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.

Now, incentive structures are often much more complicated than they are in this toy gatecrashing example. In order to show that such complexities do not undermine our reliance on Sanchirico's incentive-based story, let us quickly go through some of the ways in which things may be more complicated.

For one thing, even if we accept statistical evidence and are willing to rely on it unqualifiedly, there is still always also the possibility of individual evidence becoming available as well. And this kind of evidence may have a 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 the 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 the economic deliberation of organizations and questions about what precautions to take, how long to let the drivers drive, what level of activity is optimal, and so on. 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 point below).

The incentive structure is more complicated in important ways, even in gatecrasher-like cases, once the offender is a repeat player: just as Sanchirico's observations show that accepting character evidence will decrease John's incentives not to gatecrash if he already has an established gatecrasher record, it also shows that accepting character evidence *increases* John's incentive not to establish such a record, that is, not to gatecrash for the first couple of times. We suspect that a similar point can be made—at least sometimes—regarding the application of Sanchirico's observations to statistical evidence.

Furthermore, sometimes (in cases that are unlike the gatecrashers' case) the deliberating would-be perpetrator knows that if he chooses not

to commit the crime, no crime will be committed at all, and so the availability of statistical evidence against him will not be relevant.³³ For this reason, perhaps the incentive structure will differ between different crimes (those where typically if the agent decides not to commit the crime no crime will be committed, and those—like gatecrashing—where the crime will be committed either way and the agent’s decision only affects the identity and number of perpetrators), will differ between perpetrators with different beliefs and knowledge (for instance, about what is likely to happen if they do and do not commit the crime), and perhaps more generally will differ in a way that depends on the extent to which pieces of statistical evidence are vulnerable to strategic abuse.³⁴ And Sanchirico’s argument also seems to apply to evidence about motive as well,³⁵ but we do not typically think that relying on such evidence is that problematic.

We concede, of course, that a full and detailed discussion of the relevant incentive structure would have to address such concerns in detail (and we do just that in another article, one that also addresses doctrinal issues in detail).³⁶ But for our purposes—vindicating at least some initial difference between the appropriate legal attitude to individual and statistical evidence—not much more needs to be said here. The lesson we are taking from Sanchirico is very modest: 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 that the law gives less efficient and accurate than they would otherwise be. 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 instrumental support for the distinction we were out to vindicate.

33. We thank Mitch Berman for making us see this point.

34. In the next subsection, we are going to argue for systematic parallels between the incentive story and the epistemological story. Still, that there may be *some* discrepancies between the incentive story and the counterfactual Sensitivity story is to be expected, if nothing else because incentives are evaluated *ex ante*, and the counterfactuals *ex post*. We discuss these complications elsewhere.

35. We thank Mike Redmayne and Christian Piller for discussions relevant to this point.

36. “Statistical Evidence: The Case for Sensitivity” (unpublished manuscript).

C. 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 nonlegal cases where talk of incentives seems out of place. 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 do not, they will not.” And typically, when at a point in time some such conditionals are true, at a later point in time (some of) the very same facts are captured by counterfactuals, or subjunctive conditionals. Suppose that John proceeds to crash the gates. Then his conditional “if I do not crash the gates, they will not punish me” captures the fact that we can now—say, when John is on trial—capture with “had he not crashed the gates, we would not have punished him.” And this counterfactual should sound familiar to you: *it is the relevant instance of Sensitivity*. In other words, though the epistemological story is not itself of legal value, and though the instrumental story that is of legal value is not itself epistemologically respectable, both of them still stem from the same source, Sensitivity-style counterfactuals. These are needed both for knowledge (or are in some other closely related way epistemically relevant) and for a reasonably efficient incentive structure. While neither the epistemological story nor the instrumental one depends on the other, they are not totally independent of each other either: both of them depend on Sensitivity and related counterfactuals.

Now, 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 cannot deny that it would be *especially* nice if such a deeper story were to be found.

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 do 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 that.³⁷ 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.³⁸

How general is our emerging account? Notice that within the law, when it comes to policy recommendations, what on our account 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 that 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

37. Our focus on Sanchirico's account neither entails nor presupposes 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.

38. Wait: is DNA evidence not an obvious counterexample to our claims, a case where the law is happy relying on insensitive evidence? Things are complicated here, and we address these matters in detail in another article. Let us just note that actually the account developed here promises to explain the DNA exception to the usual suspicion with which statistical evidence is viewed. This is so, first, because DNA evidence *does* seem to be sensitive (had he not done it, in all likelihood we would have not found his DNA on the scene, and so in all likelihood we would have not convicted him), and second, because the incentive problem for statistical evidence does not seem to apply (because would-be perpetrators are hardly ever in a position to know whether DNA obtained on the scene will match theirs). But, to repeat, all this merits much more discussion, which we engage elsewhere. See "Statistical Evidence: The Case for Sensitivity."

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 or the instrumental value of the simplicity of the relevant legal rules) against relying on (insensitive) 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. In those cases, our account will not support taking statistical evidence any less seriously than we should take individual evidence. Perhaps in those cases too this would still be somewhat counterintuitive. But of course, in those cases we will have a ready explanation of the remaining intuitive discomfort, for in those cases, too, statistical evidence, while legally kosher, is not epistemically fully kosher. This, perhaps, is why we intuitively view it as suspicious. It is just that we know that the law should not place a value on knowledge or on epistemology in general (or if it does, it should place on it a value that is lexically inferior to that of accuracy), and so whatever remaining intuitions here rebel against our account should be debunked.³⁹

Are there any other nonlegal contexts to which all of this applies? Does it apply to the use of statistical evidence in shaping public policy? In allocating resources? In making personal decisions? In some contexts, relying on statistical evidence does not even give rise to an initial suspicion: some resource-allocation decisions may be of this kind, as well as perhaps decisions about taking certain precautions to prevent harm and the like. Of course, in contexts in which statistical evidence is not even initially more suspicious than individual evidence, nothing like our account is called for. If there are some such contexts where the incentive story above seems to apply, this will pose an initial challenge for our account, something that at the very least calls for explanation.⁴⁰ In other contexts, though statistical evidence is initially suspicious, epistemology

39. And there may be other debunking explanations of these intuitions as well. For one attempt at a cognitive psychological one, see K. J. Heller, "The Cognitive Psychology of Circumstantial Evidence," *Michigan Law Review* 105 (2006): 241–306.

40. We thank an anonymous referee for noting this.

matters, and so the epistemological story all by itself explains the relevant suspicion. Perhaps—we are not sure—the case of relying on statistical evidence in morality (when it comes to blaming or to resenting) is of this kind. If there are other nonlegal cases where statistical evidence is suspicious, but epistemology does not seem to matter (enough), we hypothesize that something along the lines of the combined epistemological-incentive story applies. But of course, whether this is so is something that will have to be determined piecemeal, and we hope to do that in future work.

VI. CONCLUSION

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. Perhaps the intuitive, initial, unreflective suspicion is motivated by the epistemological worries that are also exemplified by the lottery cases. But these concerns, we have argued, do not survive reflection about what considerations should guide the law. What should guide the law (perhaps among other things) are incentives, and we have found out that these too support the initial suspicion. And—as we have emphasized—the fact that the incentive-based practical considerations and the purely theoretical epistemic ones go hand in hand so nicely is not a mere coincidence: Sensitivity-like counterfactuals do the work in both.

But to an extent, the proof of legal-theoretic pudding is also in the eating. And so to complete our account, we would need to conduct a detailed discussion of the doctrines of evidence law. We would have to see to what extent our analysis (in epistemological terms and, more importantly, in incentive-based terms) fits the deeper features of these doctrines and where they call for reform. And when according to our theory they do call for reform, we would have to see whether on reflection it seems plausible that the doctrines should indeed be so reformed. One of the reasons this would be a difficult task is that legal doctrine—here as almost everywhere else—is diverse and complicated. For instance, it would be just false to say that the law today rejects all cases of statistical evidence. Things are much more complicated and nuanced than that: think, for instance, about possible distinctions between accepting statistical evidence for the prosecution or for the defense, in

criminal or civil cases, all on its own or together with other evidence, below or over a certain probabilistic threshold, and so on. As is obvious, we cannot conduct such detailed legal discussion here: a sequel to this article consists of this discussion. But we want to note both how promising the solution suggested here to the puzzle of the distinction between statistical and individual evidence is in facilitating such discussion and that the ultimate evaluation of the success of our solution will partly depend on the details of that account.