

Sensitivity, Safety, and the Law: A Reply to Pardo

David Enoch and Levi Spectre*

A belief or judgment is *sensitive* if and only if, roughly, had it been false, the relevant person would not have believed it. A belief or judgment is *safe* if and only if, roughly, the relevant person would not have held it without it being true. Both conditions – Sensitivity and Safety – seem to capture something important, at least epistemically and perhaps also legally. The best way to see this is to consider Sensitivity- and Safety-*failures*, and see how disturbing they are.

Usually, forming the belief that the table in front of you is red based on just perception seems perfectly fine. But if we find out that the table would have looked red to you even had it been some other color (say, because of lighting conditions in the room) – or, in possible worlds lingo (of which a little more below), if we find out that in the closest world (or enough of the closest possible worlds) in which it's not red you still believe it is – this seems to count against the evidential credentials of your belief. In this example, Insensitivity seems to undermine knowledge.¹

Usually, forming a belief about the time based on a quick look at a clock seems fine. But if we then find out that in circumstances not too radically different from the actual ones that clock would have shown the wrong time – or, in possible worlds lingo, if we find out that in sufficiently close worlds in which you still rely on the clock your belief about the time is false – this seems to undermine your knowledge of the time. In this example, Unsafety seems to undermine knowledge (though in this example Insensitivity may also be doing work).

* For helpful discussions and comments on earlier versions we thank Tali Fisher and two anonymous readers for *Legal Theory*.

¹ This knowledge impediment might continue to hold even if the worlds where the table is not red — say, because unbeknownst to us there is a law forbidding other colored tables under red lighting — are far enough from the actual world to render your belief Safe. If so, this belief would be safe but insensitive.

These were non-legal examples, but legal examples are not harder to construct. Thus, a factual finding is at the very least rendered problematic if it is insensitive² – that is, if, even though the finding is true, the court would have still found as it did, even had it been false. And a legal procedure that can too easily lead us astray – that is unsafe – seems at the very least initially suspicious, precisely because even if it actually leads to a true conclusion, if the circumstances had been slightly different it would have led to a false one.³

In his recent “Safety vs. Sensitivity”⁴, however, Michael Pardo argues that Sensitivity in fact plays no role in guiding or explaining legal doctrine, nor should it. En route to this conclusion, he also criticizes our attempt to utilize Sensitivity in order to shed light on the problem of statistical evidence⁵.

This paper is, to a considerable extent, a reply to Pardo. But it is, we hope, also of wider interest. Pardo’s discussion – of our previous work, and more generally – serves also as a good opportunity to revisit issues at the intersection of law and epistemology, both in specific terms (like the legal role, if any, of Sensitivity and Safety, or what to say of statistical evidence), and more generally (about how, if at all, legal theory should incorporate epistemological discussions). We happily take advantage of this opportunity, then, offering both specific critiques of and replies to Pardo’s discussion, and more general points that are of interest also independently of our exchange.

In section 1 we detail our view about Sensitivity and statistical evidence; then we defend both an epistemological and a legal role for Sensitivity (in section 2); then, in section 3, we discuss Safety, offering a

² Throughout, we apply Sensitivity and Safety to beliefs, findings, evidence, verdicts, and so on. In some contexts it may be important to be stricter on what it is that Sensitivity and Safety are meant to apply to, but not in ours, we think: We think that the most basic cases is that of sensitive (or insensitive) and safe (or unsafe) *evidence*. But all the other cases can be understood in the natural way – a belief is sensitive iff it is based on sensitive evidence, and so on. We acknowledge that this “natural way” may not be simple, and that complication may arise, but not ones, we think, that are crucial for our discussion here.

³ Note that both Safety and Sensitivity are factive conditions: If your belief is safe or sensitive, it is true. Otherwise the closest world in which your belief is false is the actual world (and presumably it is also false in most of the closest worlds where you believe that it is true).

⁴ Pardo (2018). All page references are to this paper, unless otherwise stated.

⁵ See David Enoch, Levi Spectre, and Talia Fisher, *Statistical Evidence, Sensitivity, and the Legal Value of Knowledge*, *PHILOSOPHY AND PUBLIC AFFAIRS* 40 (3):197-224 (2012). For more details engaging legal doctrines, see David Enoch and Talia Fisher *Sense and Sensitivity: Epistemic and Instrumental Approaches to Statistical Evidence*, *THE STANFORD LAW REVIEW* 67, 557-611 (2015).

more pessimistic view of its role than Pardo hopes for; finally, in section 4, we comment on the relations between epistemology and the law. If we are right, these relations are in no way straightforward.

1. Sensitivity and Statistical Evidence: What We Say, and What We Don't

If a bus causes harm, and there's a 70-percent-reliable eye-witness testifying it was a blue bus, the law has no problem finding the Blue Bus Company liable in torts. If, however, there is no eyewitness, but it's uncontroversial that the Blue Bus Company controls 70 percent of the market in the relevant area, the law (across jurisdictions, as far as we know) refuses to find the Blue Bus Company liable. The problem of statistical evidence is to explain – and if possible, vindicate – this distinction, in this and other cases. And it turns out to be no easy task⁶.

In our *Statistical Evidence, Sensitivity, and the Legal Value of Knowledge* (2012), we focus attention on the fact that while the eyewitness testimony is – if the eyewitness is reasonably good – sensitive, the market share evidence, and perhaps statistical evidence more generally, is not. Thus, had the harm been caused by a red bus, the eye witness would have probably not testified it was a blue bus, and so in the case where the court bases its ruling on the eye witness, it would not have found the Blue Bus Company liable. But the market share controlled by the Blue Bus Company would have been the same even had the harm been caused by a red bus, so a court that is willing to base a ruling on the statistical evidence would have found the Blue Bus Company liable even had it not been. Tying the discussion to general discussions of Sensitivity in epistemology (for instance, of one knowledge version of the Lottery Paradox⁷), we claim that failures of Sensitivity are often epistemically relevant – it is an epistemic bad-making feature of a belief that it is insensitive. And so, we offer the story in terms of Sensitivity as a *partial* explanation of the suspicion

⁶ For a survey of attempts, see Enoch and Fisher, *supra* note 5, section 1. For more recent attempts, see, for instance, Georgi Gardiner, *Legal Burdens of Proof and Statistical Evidence*, in James Chase & David Coady (eds.), *ROUTLEDGE HANDBOOK OF APPLIED EPISTEMOLOGY*; Martin Smith, *When Does Evidence Suffice for Conviction?* *MIND* 127 (508):1193-1218 (2018).

⁷ John Hawthorne's *KNOWLEDGE AND LOTTERIES* (2004) has many Lottery type puzzles and paradoxes. The particular one we focus on was first presented in Stewart Cohen, *Contextualist solutions to epistemological problems: Scepticism, Gettier, and the lottery*, *AUSTRALASIAN JOURNAL OF PHILOSOPHY* 76 (2):289 – 306 (1998).

towards statistical evidence. It is partial because Sensitivity only explains the intuitive reaction towards statistical evidence in central cases. Why this limited role we assign to Sensitivity?

The Sensitivity story falls short of vindicating the differential treatment of statistical and individual evidence, we argue, because the epistemic credentials of a belief or a judgment or a finding should not in themselves matter legally. In particular, and because the law's practical advantages all depend on its reliability, the law should not be willing to pay a price in reliability (or accuracy) – however small – just in order to secure some desired epistemic status (like knowledge, or epistemic justification) for its findings. More can be said here, of course⁸, and indeed we do say more below, in section 4 – but anyway, our interim conclusion here is that reliability trumps epistemological legal value. Indeed, we suspect that epistemological properties have no legal value, but we insist that if they do, this value must be lexically inferior to accuracy.⁹ (Enoch et al 2012, 213)

The bar for resolving the legal puzzle is, then, raised. It isn't enough to point out that statistical evidence has a bad epistemological profile (i.e. that it is insensitive), or unable to support knowledge for some other reason (e.g. that it is unsafe). To avoid what we call *knowledge fetishism*, an account must offer enough by way of other values in order to overcome any loss the law might incur in reliability – the proportion of correct to overall verdicts (and in the practical payoffs reliability secures)— by discounting statistical evidence. In order to vindicate the law's very different attitude towards statistical and individual evidence, then, we offer a story in terms of incentives for primary behavior (surely, the law should care a lot about *those*). The idea, in a nutshell, is that whatever incentives the substantive law is supposed to generate, they are seriously compromised unless the agent knows, ex ante, that her action will affect how the law treats her; that is, unless she has very good reason to expect the law's treatment of her to be sensitive to her behavior.

⁸ See, for instance, the thought experiment of having to choose the legal evidence regime under which your children will live (Enoch et al, supra note 5, 213). We do not have a proof that in all possible circumstances a system that is hospitable to statistical evidence is more reliable than one that isn't. Still, in many circumstances this would be the case, and anyway, those who argue for the legal significance of some epistemological notion, should be willing to defend it in face of diminished reliability.

⁹ Enoch et al, supra note 5, at 213.

The role we allocate to in/sensitivity, then, is rather limited. It has to account for our intuition that in major epistemological and legal cases statistical information is inferior. It is important to mention that our incentive-based explanation would stand even if it were somehow shown that, e.g. Safety, or Normalcy¹⁰, better match the legal or epistemological cases (a point to which we return below). In fact, putting our claim regarding knowledge fetishism aside, our account is consistent with many if not all of the epistemological accounts of the statistical evidence legal puzzle. That admitting statistical evidence would diminish the incentive to comply with the law is not challenged by other, non-Sensitivity based epistemological stories about its defective nature – though the Sensitivity-story does coincide especially nicely with the structure of explanations in terms of incentives¹¹. What we are primarily committed to is that treating statistical evidence as the law treats individual, direct evidence would diminish the incentives supplied by substantive law in sufficiently many central cases¹². Sufficient, that is, to offset the price that the law pays in reliability, if it excludes (or significantly downgrades) statistical evidence.

These are the major contours of our account. So when Pardo says that we “contend that sensitivity explains a distinction in law between statistical and “individualized” evidence.” (13), what he says may be misleading. We do not, in any way, offer explanations in terms of Sensitivity as definitional of statistical evidence. Nor do we offer any other definition – a point we return to below. Nor do we offer thoughts about Sensitivity as a vindication of the law’s differential treatment of statistical and individual evidence – as we have seen, and as Pardo notes (14, footnote 64), we argue explicitly against this suggestion, and offer the incentive-based story in its stead.

¹⁰ For the Normic account of the puzzle of statistical evidence, see Smith (2016, 2018). We briefly discuss Smith’s account in our Enoch et al (supra note 5), where we ultimately reject it for the same reason we ultimately reject a Sensitivity based account.

¹¹ Sensitivity also has the advantage – because of its structural similarity with the incentive story we give – of accounting for the more favorable attitudes the courts have toward “cold hit” DNA evidence. This is an advantage that is not mirrored, to our knowledge, by any of the other epistemological accounts.

¹² We do not want to pretend that legal doctrine is cleaner or more coherent – or less messy – than it actually is. While often (and across many jurisdictions) naked statistical evidence is not even admissible, in some cases it is, and the suspicion towards it is expressed in terms of weight or some such. For some detailed discussion of legal doctrine here, see Enoch and Fisher (supra, note 5). All that’s needed for the incentive story in the text to go through is that other things being equal, the more weight is accorded to statistical evidence the more it weakens the relevant primary behavior incentives supplied by the substantive law.

Also, we are careful to avoid too strong an epistemological commitment to Sensitivity – say, taking it to be, as some have thought¹³, a necessary condition for knowledge. We note (2012, 202) that there are problems with Sensitivity, and explicitly restrict our commitment to “something in its vicinity” being, other things being equal, an epistemological desideratum. The use we do put Sensitivity to is in explaining intuitive judgments about some central cases, in law (like Blue Bus) and in epistemology (like some lottery cases), as well as what seems like a close connection between all of these cases (across the legal and the epistemological domains). Sensitivity seems to do an especially good and elegant job of explaining these explananda¹⁴. This, we say, suffices for our discussion of statistical evidence, and is extremely plausible. So some of Pardo’s general critiques of Sensitivity, echoing the epistemological literature on Sensitivity that we cite (and to which we briefly return below), fail to engage our use thereof¹⁵.

2. Sensitivity and the Law

We find it useful to divide the discussion of Pardo’s complaints against Sensitivity into those relevant to the relation between Sensitivity and statistical evidence (in section 2.1) and the more general claims against the legal relevance of a Sensitivity requirement (section 2.2). In this more general context we also discuss the relation between Sensitivity and reliability; Pardo’s purported counterexamples to Sensitivity (in the legal context), and the question whether – objections to the legal relevance of Sensitivity aside – a positive case for its relevance can be presented.

2.1. Sensitivity and Statistical Evidence Again

¹³ Fred Dretske, *Conclusive Reasons*, AUSTRALASIAN JOURNAL OF PHILOSOPHY, 49, pp. 1-22 (1971); Robert Nozick, PHILOSOPHICAL EXPLANATIONS, (Cambridge, Mass.: Harvard University Press), 1981; Keith DeRose, *Solving the Skeptical Problem*, PHILOSOPHICAL REVIEW 104:1, 1-52 (1995). DeRose imbeds Sensitivity within a contextual account the result of which is an account that differs in important ways from the former two.

¹⁴ In “Statistical Resentment” (manuscript) we extend the discussion even further, to cover also cases of statistical resentment, and many others. The results for Sensitivity, and indeed, for our incentive-based story, are mixed.

¹⁵ The extent to which this is a criticism of Pardo is mitigated by the fact that for him statistical evidence – and certainly, *our* discussion of statistical evidence – is not the main theme. Even if these points fail to engage our discussion of statistical evidence, they may still be relevant to his general discussion of the legal (in)significance of Sensitivity.

Pardo's main complaint against the relevance of Sensitivity to an understanding of the problem of statistical evidence is that the two distinctions – sensitive/insensitive, statistical/individual – cut across each other. But his discussion here suffers from too broad an understanding of statistical evidence.

When we – following the literature – speak of statistical evidence, we think of examples such as Blue Bus, and the phenomenon it is an example of. This is the phenomenon sometimes called base-rate evidence, sometimes market-share evidence, or *naked* statistical evidence. This is the phenomenon we offer thoughts about Sensitivity as relevant to (though in the somewhat restricted, tentative sort of way described in the previous section). But when Pardo speaks of statistical evidence, he seems to be thinking of any evidence at all that involves a statistic. We never claimed – who would? – that any evidence involving statistics is poor or should be inadmissible, or indeed is insensitive.¹⁶ So Pardo's example (15) of *Tyson Foods*¹⁷ - where statistics were admitted in order to prove something about the circumstances, rather than as direct base-rate data for guilt or liability – is just irrelevant for our discussion. The fact that *such* evidence may be sensitive is thus neither here nor there.

How *do* we, then, define statistical evidence? We don't. We – again, in a way that's consistent with the theoretical literature on statistical evidence – start with the examples. They clearly capture something intuitively important. We then try to understand the relevant phenomena better. If we're fortunate, we may *end up* with a definition, or an analysis¹⁸. Or we may not – but may still gain understanding. In such a case we may need to settle for an "I can't define it, but I know it when I see it" attitude¹⁹. Depending on the expectations you started with, this may be frustrating. But even if frustrating, this would be better than settling for a delineation of "statistical evidence" that clearly misses the underlying intuitive concerns. Grouping under this title all evidence having to do anything at all with statistics – as Pardo seems to do – is guilty of this much more serious flaw.

¹⁶ See, for instance, our discussion of DNA evidence, especially in Enoch and Fisher (2015). In some cases, we argue, DNA evidence may be both statistical and good evidence.

¹⁷ *Tyson Foods, Inc. v. Bouaphakeo*, 136 S. Ct. 1036 (2016).

¹⁸ We are here echoing Kant on the different place definitions have in mathematics and in philosophy. See Lewis W. Beck, *Kant's Theory of Definition*, *THE PHILOSOPHICAL REVIEW* 65, 179-191 (1956, 187-8) and the references there.

¹⁹ Paraphrasing Justice Potter Stewart's famous phrase about the test for obscenity in *Jacobellis v. Ohio* (1964).

2.2. Sensitivity and the Law More Generally

Of course, if Sensitivity sheds light on the phenomenon of statistical evidence (properly understood), this suffices to show that it is relevant to the law of evidence, thus falsifying Pardo's more general claim to the contrary. But putting the specific discussion of statistical evidence to one side for now, we want to engage Pardo's more general arguments against the legal relevance of Sensitivity. We find here general epistemological doubts about Sensitivity, thoughts about the relation between Sensitivity and reliability, and specific purported counterexamples. We address these in turn. We conclude this section with a hint at a positive case for Sensitivity.

2.2.1. General doubts about Sensitivity

Epistemologically, the main general concern Pardo cites about Sensitivity is that a mistake in a sufficiently distant possible world needn't matter. Thus, if I believe that *p*, and my belief is insensitive, this means that in the nearest neighborhood of possible worlds in which *p* is false, I still ("most probably" we would add) believe *p*.²⁰ But this, argues Pardo (following an influential line of epistemologists), need not compromise my belief's epistemic status, *if those worlds – the closest in which p is false – are sufficiently distant*. Sosa's Trash Bag case is a powerful example²¹: If I throw a trash bag down the chute, and believe (after a short while) it's made it to the basement, my belief is not rendered any less justified (and is not precluded from amounting to knowledge) just because it is insensitive: For had bag not made it to the basement, this would have been because it was snagged somehow in the chute, and then I would have still believed it's made it to the basement. Presumably, the insensitivity of my belief here is rendered epistemically moot because the bag not making it to the basement is such a rare occurrence and distant possibility (11).

We agree – and as noted, were careful to say so explicitly – that things are complicated in the vicinity of Sensitivity, and that Sensitivity may not be necessary for knowledge. So we can accept such examples at face value. But more can be said.

²⁰ See Nozick's (*supra*, note 12) discussion of why neighborhoods of possible worlds deliver a better heuristic for determining whether a belief is sensitive (or whether if the agent were to believe that *p*, *p* would be true – his fourth, Safety-like condition).

²¹ See Ernest Sosa, *How to defeat opposition to Moore*, *PHILOSOPHICAL PERSPECTIVES* 13, 137-49 (1999).

First, of course, this is just one kind of example. From the fact that Sensitivity doesn't seem especially important here it doesn't follow that it isn't elsewhere. Seeing that we have already given up on the hope for full, exceptionless generality, our claim – that Sensitivity is very often epistemically relevant (especially in capturing cases where a true and highly probable belief *does not* amount to knowledge) – is not refutable by counterexample. Of course, this does not mean it is immune to challenge and counterargument. The claims that smoking causes cancer and that symmetry is often a beauty-making feature of a face are not refutable by single counterexamples, but are not meaningless, and may be challenged in other ways. If a case is to be made against the (attenuated) epistemic significance of Sensitivity, then, it must be made not by a clever counterexample, but by going deeper – showing that tracking the truth is not after all important epistemically, say, or by showing that Sensitivity fails in a large subset of paradigmatic, central cases, and so on. We don't think that such objections to Sensitivity are to be found in the literature.²²

Second — and this is something we discuss a bit more when we turn to Pardo's purported counterexamples below — the belief in the Trash Bag case may be technically insensitive, but it still falls into the incentive pattern that is more pertinent for our account. If the relevant agent is faced with a decision whether or not to illegally throw the trash bag in the basement container, she knows that her decision dramatically increases the likelihood of her being punished. So it's not as if this case counts against the legal significance of the incentive-version of our Sensitivity story²³.

Third, we confess that we don't think a better, watertight formulation of Sensitivity – one that will render it necessary for knowledge – is forthcoming.²⁴ The problem is not so much that there's no way of

²² We thank both anonymous referees for pressing us on related points.

²³ This does not mean, as one referee suggested, that we're giving up on the relevance of Sensitivity, proceeding to do things in terms of (Sensitivity-unrelated) incentives alone. First, although in atypical cases the incentive story and the Sensitivity story come apart, still in the more paradigmatic cases they do not. Second, although the belief that the trash bag is in the basement is not in this case sensitive, beliefs in its vicinity are (say, that I *tried* to get the trash bag down to the basement), and such beliefs are relevant here both intuitively and for the incentive story.

²⁴ Additional well-known counterexamples to Sensitivity being necessary for knowledge are Kripke's red fake barn (Saul Kripke, *Nozick on Knowledge*, in *PHILOSOPHICAL TROUBLES. COLLECTED PAPERS VOL I.* Oxford University Press (2011), 162-224)and Williamson's brain in a vat climber (Timothy Williamson, *KNOWLEDGE AND ITS LIMITS*, Oxford University Press (2000). To take Williamson's example: though I know I am not a brain in a vat who believes I'm climbing the Himalayas (as it doesn't even seem to me that I'm climbing the Himalayas), this belief is insensitive. In the closest

addressing the problem above. Rather, it's that any amendment is expected to render Sensitivity problematic in other respects²⁵.

Also, it's important to keep in mind that Sensitivity is advanced – at most – as a necessary condition on knowledge (and let us remind you that we go only for even weaker claims). No one, as far as we know, thinks of Sensitivity as sufficient for knowledge. So even if no insensitive belief can be knowledge, this does not mean that all sensitive beliefs constitute knowledge. In fact, thinking of Sensitivity as necessary for knowledge is consistent with insisting also on the necessity of Safety. Indeed, Nozick — an early Sensitivity advocate (supra, note 12) — incorporates a Safety-like condition as part of his Tracking Theory of knowledge. And Williamson²⁶ (2000, 2009) advances Safety as a (non-reductive) sufficient condition for knowledge. Needless to say, both Nozick's theory and Williamson's have epistemological and theoretical advantages. Pardo may worry that Sensitivity is not while Safety is a candidate for a sufficient condition for knowledge. But this is not a flaw in Sensitivity, or a reason to think it is epistemically (or legally) irrelevant. To the extent that Pardo's criticism is motivated by the thought that Sensitivity is not sufficiently close to knowledge, then, it misses its mark.

Much more, then, remains to be said about Sensitivity, both in "pure" epistemology and in applying it elsewhere (for instance, in a legal context). For the limited purposes to which we enlist Sensitivity, though, much of this is not needed, and so we can leave much of this work for another day (though some related issues will resurface as we discuss Pardo's purported counterexamples).

worlds where that belief is false – that is, where I am a brain in a vat who believes he's climbing the Himalayas – I still have that belief (namely, that I'm not a brain in a vat who believes he's climbing the Himalayas).

²⁵ One major issue is whether in applying Sensitivity to a belief we're holding the evidence fixed or not. In Sosa's example we arguably only get the result that the belief is insensitive because we're holding the agent's evidence – that he's thrown the bag down the chute – fixed. Otherwise, the closest world where his belief is false is one where he doesn't even throw the bag, in which case nor does he believe that the bag is in the basement, and the belief *is* sensitive after all. We could play with the relevant belief – that the bag is in the basement? That it's *reached* the basement? – but this looks like cheating. Whichever way we go on whether or not the evidence should be fixed seems to call for justification, and indeed to raise problems. We return below to the question what is held fixed in checking for Sensitivity.

²⁶ Williamson, supra note 23. See also Timothy Williamson, *Replies to Critics*, in: Patrick Greenough, Duncan Pritchard & Timothy Williamson (eds.). WILLIAMSON ON KNOWLEDGE. Oxford University Press, 280-384 (2009).

2.2.2 Reliability

Pardo repeatedly complains that Sensitivity fails to track reliability, or reliability of the evidence in nearby possible worlds. Evaluating this claim is rendered difficult by the fact that he nowhere explains what he means by reliability, and that it seems he understands it in a somewhat unique way.²⁷

Often, reliability is understood as frequency of truth, or truth propensity: Namely, a belief-forming method is reliable if and only if it tends to lead to true beliefs, if the ratio of true beliefs is high enough among the beliefs it generates, or perhaps if this ratio is high enough in the relevant set of cases or circumstances. A thermometer, for instance, is reliable if its readings are accurate (or accurate enough) often enough, in non-extreme temperature and pressure conditions. If this is how reliability is understood, then *of course* Sensitivity and reliability come apart²⁸ in the following sense: Reliable beliefs can fail to be sensitive (or safe), e.g., a belief that one will not win a large lottery. But this claim that Sensitivity and reliability thus understood come apart is one that we – as Pardo knows – wholeheartedly endorse. Surely this is a feature of our view, not a bug. After all, beliefs (and legal judgments) formed on the basis of sufficiently strong statistical evidence can be *extremely* reliable (and much more reliable than many formed on the basis of eyewitness testimony, for instance).

Perhaps, then, Pardo means more by “reliability”. Perhaps he means reliability in a way that already incorporates something modal (say, *non-accidental* frequency of truth), or perhaps this is why he often speaks of reliability *in nearby possible worlds*, perhaps allowing that a sensitive belief is reliable in the actual world, but perhaps not in nearby worlds? Now, it’s true that Sensitivity comes apart from such modal conditions as well, and that one way of understanding Safety renders it immune to a similar accusation: On one way of understanding Safety (and one way of understanding reliability), Safety *just is*

²⁷ See Pardo’s footnote 48 for an attempt – one whose details are not clear to us – to explain what is meant by “reliability”. Also consider: “The reliability of evidence is distinct from whether a belief is sensitive. The reason for this—and this reason will be an important reason when assessing legal evidence—is that the closest possible world in which a belief is false may not be a close possibility.” (11) The first sentence here is true but unproblematic, as we are about to argue. What follows it is unclear: there is, as far as we can tell, no straightforward connection between the modal property Pardo describes and reliability.

²⁸ Pardo speaks of Sensitivity (or Safety) *tracking* reliability, but he can’t mean by that more than just a claim about them not coming apart too often. So this is what we focus on in the text.

reliability (say, of the relevant belief-forming method) in close possible worlds. But if this is how Pardo understands the relevant reliability requirement, then accusing Sensitivity of failing to satisfy it begs the question against it. It may amount to just accusing Sensitivity of not being Safety²⁹ – a true accusation, but one lacking dialectical force in our context. What Pardo needs is a reason to think that reliability *thus understood* is an important desideratum, legally and epistemically. He hasn't given such a reason, nor do we think he can.³⁰

Before concluding this section, we want to address Pardo's hope (e.g. 20) that Safety does better than Sensitivity when it comes to reliability. We have already highlighted worries about how to understand reliability (and whether there's a serious problem here for Sensitivity). Furthermore, on any non-tendentious (or circular) understanding of reliability, Safety and reliability come apart as well. But we now want to emphasize the following point: If the relevant game, for Pardo, is all about reliability, then why not just require reliability, and get it over with? If Safety is better than Sensitivity, according to Pardo, because it scores better on reliability (however precisely it is understood), then why not cut the middle man, and just go straight for reliability? This is not a move we recommend, of course, partly because statistical evidence can be extremely reliable even when clearly and uncontroversially suspicious. But given what Pardo says, it is hard to see how he can avoid this move.

Of course, we agree that some version of a reliability requirement is extremely important legally³¹. But we have not yet been given any reason to believe that the reliability requirement that *is* significant is one that Sensitivity can violate and even if it can, how it could do so often enough.

²⁹ "Second, sensitivity does not track the risk of drawing erroneous inferences from evidence in close possible worlds." (10)

³⁰ There is reason to suspect that even if Pardo means to be working with a reliability-supporting Safety condition, this intention must ultimately fail. The notion of Safety Pardo is working with is Sosa's (supra, note 20)— e.g., in his use of Sosa's trash bag example (p 146). This construal of Safety doesn't square well with Pardo's reliability claim. In response to the worry that modal conditions are not deductively closed, Sosa says: "Yes, in fact this is one reason why our account of Safety is only a first approximation. Here now is a closer approximation. What is required for a belief to be safe is not just that it would be held only if true, but rather that it be based on a reliable indication." (p 149). In other words, Safety certainly doesn't secure reliability, reliability must be added to it externally as an extra condition.

³¹ Pardo often accuses Sensitivity of being a poor guide to *probative value*, but it is unclear what he means by that phrase. Sometimes it seems that he just takes high probative value to be the property of good evidence. If so, accusing Sensitivity of having poor probative value is of course a serious accusation, but it needs substantiating (it's not *the substantiation*). At other times he seems to think of probative value as very closely connected to reliability. In

2.2.3 Counterexamples

Pardo also brings purported counterexamples to the legal relevance of Sensitivity, and we think it will be helpful to discuss two of them at length.

Consider, then, the example Pardo calls *Drug Weight*:

A criminal defendant is tried for possessing a large amount of an illegal substance. According to the criminal statute at issue, the prosecution must prove that the defendant possessed 500 grams or more of the substance. A chemist who sampled and tested the substance will testify as an expert that the amount seized was over a kilogram. The expert is well qualified, employed an acceptable and reliable methodology, and otherwise satisfies the criteria for admitting expert testimony. The chemist's process, however, tends to overestimate weight by a very small amount (say, by one gram or less). (17)

Pardo claims that the closest possible world in which the defendant possessed less than 500 grams of the substance is one where he possessed *almost* 500 grams of it, and in that world, the chemist (prone as he is to slight overestimation of weight) would have testified it was over 500 grams. In that case, the conviction (even in the case of the 1 kilo) is insensitive. But, Pardo claims, the chemist's testimony is – in the kilo case – perfectly fine. So Sensitivity doesn't generate the right result here.

The main problem with this example is that it's not at all clear the evidence here *is* insensitive. The reason for that is that it's not at all clear that the closest possible world in which the defendant is not guilty of possession is the one where he possesses just under 500 grams. True, it's not easy to fill in the details for a proximity relation between worlds in non-metaphorical terms, but first, it's not as if nothing can be said, and second, even without further details, it's clear that Pardo's reasoning here doesn't work. Pardo seems

that case, the discussion in this section applies. For a discussion of statistical evidence that utilizes one precise notion of probative value – that of likelihood ratio – see Marcello Di Bello, *Trial by Statistics: Is a High Probability of Guilt Enough to Convict?*, forthcoming in MIND

to think that because the closest possibility *on the weight metric* where there's no violation is just under 500 grams, this entails that this is the closest possible world. But this is just not so. In order to know what the closest possible world is in which the defendant doesn't possess over 500 grams of the drug we need to know much more about the case: We need to know, now in intuitive rather than possible-world talk, what would have happened had he not possessed so much of the drug. Perhaps had he not possessed a kilo, he wouldn't have possessed any of the drug, because only a kilo makes it worth his while to get into this game to begin with. Perhaps the drug comes in packets of 200 grams, in which case, a world in which he possesses just under 500 grams is very distant (compared to a world in which he possesses 600 grams or 400 grams). If this is how the details of the case are to be filled in, then the chemist's testimony is straightforwardly sensitive, and all is well. Well then, can't Pardo insist on filling in the details in such a way that makes the 499-gram world the closest one in which the defendant doesn't possess over 500 grams? Of course he can. But we need to see those details filled in. And Sensitivity predicts – quite plausibly, we think – that in *such* a scenario, there *is* something problematic with the described chemist's testimony.

Moreover, consider our primary explanation of the problem with statistical evidence, in terms of incentives, and now assume that we fill in the details as best we can in favor of Pardo's case. In order for the criminal defendant's incentive (to comply with the weight restriction) to diminish, he would have to know about the chemist's tendency to overestimate the weight of the illegal substance. Even though in that case he would have no reason to possess 499 grams rather than higher weight, he would still be incentivized to settle for, say, at most 450 grams. This is hardly a case that poses a problem for our account.

Next, consider *Lake Pollution*:

The defendant is charged with illegally dumping a toxic substance into a lake. A video clearly shows the defendant dropping a container of what appears to be the toxic substance down a long drain pipe that leads to the lake. There is a small hook in the pipe, however, and it is possible that the container could have caught on the hook and not reached the lake, although it would be extremely difficult for this to have occurred. (17)

Again, Pardo thinks the video evidence is insensitive, and fully kosher. But again, he is not careful about the precise formulation of the relevant counterfactuals, or about the possible world relevant for evaluation of Sensitivity here. What happens on the closest possible world in which the defendant is not liable? (Or perhaps – on the closest possible world in which the defendant does not dump a toxic substance into a lake?) And again, this depends on the details. Perhaps, for instance, had the defendant not dumped the substance into the lake, he would have gotten rid of it in some other way. In that case, the video would not have shown him dumping it in the lake, and so is sensitive. Of course, this is a hypothetical case, and Pardo is entitled to fill in the details in whatever way he wants. It's just that if he fills in the details in the way needed to render the video evidence insensitive, we predict that proclaiming it perfectly kosher will no longer be at all intuitively plausible. Moreover, just as in the previous case (as well as with the trash bag case on which this case is obviously modeled) there is no reason to suspect that the relevant incentive is reduced in any way by incorporating the video evidence.

Pardo's counterexamples to the legal relevance of Sensitivity are thus unconvincing. Before concluding this section, though, we want to address a complication, one that is relevant to the evaluation of Pardo's purported counterexamples, but also more generally to any application of Sensitivity (or Safety, or indeed any other counterfactual condition)³². We've been insisting that in many cases admitting statistical evidence will entail a Sensitivity violation – it will make it the case that even had the person in front of the factfinders not committed the crime or the tort, the factfinder would still have found against them. But this may have been too quick. Many crimes, for instance, are quite rare, and the closest possible world where the defendant doesn't commit the crime is a world where the crime doesn't occur at all, and ipso facto there's no trial, and no conviction. In such a case, then, a conviction based on statistical evidence is sensitive after all (it's not true that the court would have convicted even had the defendant been innocent; had the defendant been innocent, there would have been no trial). More generally, when applying Sensitivity, we need to decide what to hold fixed as a part of the antecedent: Do we ask what would have happened had the defendant been innocent? Had he been innocent but still brought to trial?

³² We thank both referees for pressing us on related points.

Had he been innocent, brought to trial, and faced similar evidence? Etc. (And as you can see, the need to decide what to hold fixed in applying a counterfactual applies to any counterfactual condition). Crucially, the answer to this question is not given by Sensitivity (or Safety) itself – it's a question that must be answered before Sensitivity is applied (or – which comes down to the same thing – it's a question about the precise version of Sensitivity to be applied).

The way to make progress on this question is, then, not to consult Sensitivity, but rather to consult the motivations and intuitions underlying it. Here's one such hint, then: We want, to the extent possible, an account that explains the suspiciousness of statistical evidence across all the cases in which it's suspicious. So an account according to which the verdict of Sensitivity depends on how common the relevant crime is can't succeed. Statistical evidence is suspicious – for the same reasons – when relevant to a rare crime (such that had the defendant not committed it no one would have) and to common ones (where, as in Gatecrashers, had the defendant not committed the crime, many others would). So if a Sensitivity-style story is to have any chance at all, the fact that the crime (or tort, etc.) has occurred must be held fixed, as a part of the counterfactual's antecedent. And this is not ad hoc – it's not a measure introduced just to save Sensitivity here. It also jibes well with the underlying concern. If you're the factfinder in a trial, what would have happened had there not been a trial (or had you not been called to jury duty, or had another judge been assigned the case, or ...) is not a relevant question for you. What seems needed is to hold fixed the occurrence of the trial, and ask whether you would have acquitted had there been a trial, and had everything else been pretty much as it actually is, except that the defendant had been innocent. Interestingly, this also coincides well with the Sensitivity-related incentive story: Presumably, the incentive we should be most concerned with here is the incentive the potential gatecrasher has when already at the stadium gates – not incentives, say, to become or not to become a football fan from the get-go, and so on.

We don't want to pretend that what we've said here suffices as an answer to the question what to hold fixed when applying Sensitivity. But enough has been said, we think, to show how, in broad outline, to

address this important complication³³. And let us again emphasize that to the extent that you find this concern troubling regarding Sensitivity, you should find the parallel concern troubling regarding Safety (and any other counterfactual condition) as well.

2.2.4 A Positive Case for Sensitivity's Legal Relevance?

Pardo's negative case, then, fails. Can anything be said, though, in favor of Sensitivity's legal relevance? We should note that we are here at a somewhat awkward dialectical spot. After all, as we emphasize in our original paper (and as Pardo notes only quickly), at the end of the day we think that *nothing* epistemological should intrinsically matter legally, Sensitivity (and Safety) included. So we offer the following thoughts very tentatively: There are reasons to think that if any epistemic properties are a legal desideratum, Sensitivity is high on the list.

One point here comes directly from Nozick (supra, note 12), who was among the first to introduce modal conditions to an analysis of knowledge. The intuition he's working with is that in order to count as knowledge, a belief must *track the truth*. This idea – of tracking – lends immediate intuitive support to a Sensitivity condition, according to Nozick the main working necessary condition on knowledge together with a Safety (-like) condition: If it were true that p, you would believe that p (via the same method). A related intuition is that in order to count as good evidence for p, a piece of evidence (or perhaps an epistemic procedure) has to differentiate between p and not-p. And again, a very natural way of fleshing out this intuition leads to Sensitivity.

³³ Let us mention another kind of case that is relevant here. Cases may be described when the availability of the statistical evidence itself may depend on the defendant's behavior. For instance, a specific gatecrasher may be a social magnet, such that his gatecrashing may make it the case that many more gatecrash. In such a case, the statistical evidence against him may be sensitive after all (because had he not gatecrashed, the statistical evidence would not have been the same, in fact, it would have been much weaker, perhaps too weak for conviction even under an evidence regime that admits naked statistical evidence). For this kind of case, see Michael Blome-Tillman, *Sensitivity, Causality, and Statistical Evidence in Courts of Law*, THOUGHT 4, 102-112 (2015), at 106-7. But either the factfinder knows that the person in front of them has this kind of social effect, or they don't. If they do, then the very statistical information – that so many people crashed the gates – is direct, non-statistical and unproblematic evidence against the defendant. And if the factfinder doesn't know about the defendant's social status, then, we want to suggest, it makes sense to hold fixed the percentage of those attending the stadium who crashed the gates, for reasons similar to the ones in the text.

These were general epistemological points, which apply in legal cases as well. But when it comes to the law, more can be said. Think about a fact-finder who is rather dogmatic about a defendant's guilt. This fact-finder's judgment may be safe – if, that is, the closest possible world in which the defendant is innocent is sufficiently distant. Even in such a case, though, there's a clear flaw in this fact-finder's attitude. And this flaw is nicely and cleanly explained by Sensitivity: we require our fact-finding procedures to not be dogmatic in such a way, that is, roughly, to be such that had the defendant been innocent, they would not have issued a guilty verdict.

More can be said, of course, about the merits of Sensitivity – both absolutely, and compared to Safety (and we say more about Safety in the next section). But we should stop here. And let us remind you that we're not trying to push Sensitivity as the single overarching epistemic virtue of evidence. We are open to the possibility that Sensitivity matters in some circumstances more than in others, that sometimes Safety is important, and so on. It's just that we reject Pardo's negative conclusions about Sensitivity's legal relevance.

3. And Now: Safety

Forget Sensitivity, then. How well does Safety do, in terms of both epistemological and legal relevance? We discuss, in section 3.1, general worries about Safety, and then, in section 3.2, specifically the application to the law, and even more specifically, to statistical evidence.

3.1 Safety, in General

Perhaps the most pressing need is to fill in the details of a Safety requirement, in a way that both maintains the connection to the intuition underlying it (about the possibility of a mistake being sufficiently distant) and avoids dangerous counterexamples³⁴. (Pardo nowhere states an official statement of Safety, or for that

³⁴ It needs to be said, though, that on the most developed construal of the Safety condition (Williamson, *supra* notes 23 and 25), we might just as well talk about knowledge. This is not only because Safety is a sufficient condition for knowledge according to Williamson, it is also because in order to determine — specifically in complex cases — whether there are worlds (possibilities of error) that are too close for a belief to be Safe, we will have to determine that by appealing to knowledge itself.

matter of Sensitivity either.) The problem is not just the inherent vagueness in thoughts about the belief remaining true in *sufficiently close* worlds – perhaps such vagueness is just something we’re going to have to live with. The problem is that it’s not clear what to say of cases where there are several equally distant worlds, such that in some of them – perhaps most of them – the relevant belief stays true, and in others it’s false.

Consider the example Pardo discusses in footnote 50 (11)³⁵. In that example, there is one liar (bad apple) among a million service people employed to answer questions about Apple products. Here, Pardo claims plausibly, if you call, randomly get one of the many honest operators, and believe her answer (p, say), your belief is not threatened by the fact that had you gotten the liar, you would have believed falsely. The first thing to note about this is that, though true, this is irrelevant to Sensitivity: The counterfactual question relevant to Sensitivity is: Had p been false, would you have still believed it? Well, in the closest possible world in which p is false it is still overwhelmingly likely that an honest operator answers your call. So she tells you that p is false, and you no longer believe p. The counterfactual Pardo considers – what would have happened had the liar picked up your call – is not the one relevant for an evaluation of Sensitivity. In fact – and more directly relevant to this section – this example is more challenging to Safety than it is to Sensitivity. For the world in which my procedure – relying on the operator who picks up my call – leads me astray is just as close as the worlds in which it doesn’t (getting the liar to take your call is unlikely, but not any closer a possibility than getting any other of the operators). Depending on how exactly Safety is spelled out, then, this may mean that the belief here – though justified (and sensitive) – is unsafe.

Perhaps somewhat more generally, let us assume that knowledge is closed under multi-premise logical operations such as conjunction, that is, that if an agent knows that p and knows that q, she knows (or at least can know) the conjunction p-and-q. Suppose further, that I count as knowing that the bag has reached the basement every time I throw it in the chute and in fact it does reach the basement. With these two assumptions, trouble ensues. By introducing more and more of these conjunctions, I can reach the

³⁵ Pardo takes this example from Ernest Sosa, *JUDGMENT & AGENCY* (2015, at 119). Interestingly, Sosa uses it to argue against Prichard’s *Safety* condition according to which for a belief to be safe, it needs to be true in *all* nearby worlds. On Sosa’s Safety condition, it need not be true in all of them.

conclusion that the trash bag has reached the basement in all of the times I threw it down the chute. Furthermore, I can *know* that this is so. But this event – the conjunctive one, about the bag reaching its destination in all of those times – is of very low probability, as I can also know. (Even if the probability of each bag reaching the basement is very high, the probability of the conjunction can be very low.) This, of course, is a problem. And it stems from the fact that if the closest worlds to the actual one in a single trash-throwing event are worlds where the trash bag gets to the basement, this modal fact (arguably) stays fixed for more and more events of a similar nature. Throwing more bags, after all, doesn't make snagging worlds closer. Seeing that probabilities do not behave in a similar way, this may be a problem for Safety (but not for Sensitivity): A belief cannot be Safe from being false and yet highly improbable.

How can a Safety theorist deal with this problem? One way (perhaps the only way) is to relax the assumption that in *all* nearby worlds p needs to be true in order for p to be Safely believed. If this is so, every trash-bag-throwing-event in N has some worlds in the vicinity of w_N where it doesn't reach the basement. In w_N you will not count, then, as knowing that *none of the N bags snagged* because there is a large majority of worlds where at least one of them is snagged. Problem solved? Not at all.

Sosa's case raises a pertinent problem for Safety. As Kvanvig puts it³⁶, if Sosa's Safety condition isn't transgressed in the trash bag example, it can only be because unlikely events don't take place in close possible worlds. And this means that we cannot appeal to the Safety construal of the previous paragraph. If we did, we would face the problem that lottery propositions – the propositions that resemble the statistical evidence cases – would count as known, a huge blow to at least one central motivation for Safety accounts, namely, to explain the absence of knowledge in those very cases. But even putting the problem from the previous paragraph to one side, it's not clear why snagging worlds (as we may call them) are not in the vicinity of the world under evaluation while worlds where you win the lottery are. If so, either in both Lottery and Trash Bag cases the agent knows (that the ticket is a losing one, and that the trash bag has reached the basement, respectively), or in both he doesn't. Either way, unlike Sensitivity, Safety gives us no

³⁶ See Jonathan Kvanvig, *THE VALUE OF KNOWLEDGE AND THE PURSUIT OF UNDERSTANDING*, Cambridge University Press (2004), at 135.

guidance here and cannot explain how one proposition is known while the other isn't. And of course, all of this is directly relevant to what Safety has to say about statistical evidence cases, as we're about to see.

Before getting to that, though, we want to emphasize one lesson that can be adduced from this and other examples. The Sensitivity condition explains the difference between lottery and the paradigmatic non-statistical cases in a clear (often striking) way. In the closest worlds where your belief *that your ticket loses* is false, your evidence is just as it is in the actual world (where you are not as lucky), and you still believe your ticket is a losing ticket. As should be evident from the discussion of the Trash Bag case, Safety has less to offer in this regard. In fact, Safety theorists themselves do not agree on whether lottery propositions are known³⁷. Of course, if the best version of a Safety theory proclaims lottery propositions known, then Safety is strikingly inadequate as an explanation of the problem of statistical evidence.

3.2 Safety and Statistical Evidence

What, if anything, does Safety have to say about paradigmatically suspicious cases of statistical evidence?

Recall that cases like *Blue Bus* are handled by Sensitivity with ease. Can Safety do as well? Here is what

Pardo says in order to convince us that the answer is positive:

“The eyewitness evidence will be safer when there are close possibilities in which similar accidents are caused by non-blue buses. In close worlds in which another company causes the accident, the witness (who, by hypothesis, has an accuracy rate of 70 percent) is unlikely to testify that the bus was blue. By contrast, in close possible worlds in which another company causes the accident, the market-share evidence remains the same. For example, suppose there are ten similar cases except that the Red Bus Company caused the accidents. Relying on the market-share evidence in those cases would result in ten errors against the Blue Bus Company and zero correct decisions. By contrast, relying on a witness who is 70 percent accurate would be expected to produce seven correct identifications (“the bus was red”) and three errors (“the bus was blue”).” (21)

³⁷ For instance, Sosa (*supra*, note 33) – on whom Pardo often relies – thinks lottery propositions may be known, and attributes the opposite impression to a confusion between a Safety condition (which he endorses) and a Sensitivity one (which he rejects).

We don't find this paragraph convincing. It's true, of course, that the market-share evidence – unlike eye-witness testimony – leads to falsehoods in close worlds where the bus that caused the harm is not a Blue Bus bus. But taking this to be relevant is at least as naturally understood as applying Sensitivity as it is applying Safety³⁸ – arguably, it amounts to noting that if we rely on the statistical evidence, then had the harm been caused by a bus belonging to a different company, we would have still found the Blue Bus company liable (Consider again: “in close possible worlds in which another company causes the accident, the market-share evidence remains the same”)³⁹. Further, Pardo's choice of an example is tendentious here: True, if we only count ten (atypical) cases in which the Red Bus Company caused the harm, then relying on the eyewitness will lead us astray less often than relying on the market-share evidence. But if we think of ten (more typical) cases in which the Blue Bus Company caused the harm, we get the opposite result – relying on the market-share evidence doesn't lead us astray at all, and relying on the eyewitness does (in three out of ten cases). And in the relevant cases – when the fact-finder doesn't know ahead of time which company is responsible for which harm – well, in those cases all we have to work with are the probabilities, and in those terms the eyewitness testimony and the market-share evidence are alike. So Pardo's attempt to show that Safety handles the statistical evidence cases well is unsuccessful.

Regardless of what Pardo says about this, how does Safety handle statistical evidence cases? The relevant question to ask about Blue Bus is the following: When we base our judgment (that it's a blue bus) on the market share evidence, and when that judgment is true, is it also true in the sufficiently close possible worlds in which we so believe? (And the analogous question for relying on the eyewitness.) The answer is not at all clear. One relevant problem is an instance of the general one from the previous subsection: For Safety to do work here, we would need a version of it that is not too strict (requiring truth in *all*

³⁸ If Piller (*Beware of Safety*, forthcoming in ANALYTIC PHILOSOPHY) is right, this is not just a one-off mistake. He claims that the appeal of Safety and the mistaken assumption that it is an epistemological property to begin with, is owed to its similarity with Sensitivity. And as already noted, Sosa (supra, note 33) argues in the opposite direction that our temptation to resort to Sensitivity confuses us into thinking that we do not know that our lottery tickets will lose (when they do).

³⁹ It's not entirely clear what to say of the eye-witness-testimony case here. Clearly if we believe that *p* based on the testimony of a 70%-reliable eyewitness, there are close worlds where we believe the witness but she is mistaken. So Safe belief (that *p*) seems out of the question here, as is, for the Safety theorist, knowledge (that *p*). Perhaps the testimony suffices for weaker results (like knowledge that probably-*p*). However, we can get such knowledge from the market-share evidence as well. Whatever epistemic merits Safety might have, then, it seems ill-suited for the law even for the most basic case. See further discussion below.

closest worlds) but that is not too weak to allow for knowledge of lottery beliefs (for such a version will also allow for knowledge based on statistical evidence). Another problem is that in order to know whether any of the red-bus worlds are closer to the actual world than all the blue-bus worlds we need more details about the case. In some ways of filling in the details – say, when the Red Bus company has much stricter safety regulations than the Blue Bus company – red-bus worlds are going to be further than blue-bus worlds, and the statistical-evidence will be safe. In others – where the risk-of-accident-per-bus is roughly equal across companies, and red buses pass in the same route often enough – red-bus-worlds will be very close, and the statistical-evidence-based judgment will be unsafe. Of course, intuitively, the statistical evidence is suspicious (though perhaps not equally suspicious) across these different scenarios. So this is already a problem for Safety (and an advantage for Sensitivity, which is not equally vulnerable).

Pardo claims that Safety fits better with legal practices. Partly for the reasons above (for instance, about statistical evidence), we doubt it. But we want to note, in concluding this section, that there needn't be a unified account here. It's quite possible that in some contexts legal practices and doctrines incorporate a concern for Sensitivity (or something like it), in others for Safety (or something in its vicinity), and in others yet other concerns. We do not see any grounds for expecting greater coherence from the law (or, for that matter, even from epistemology).

4. On the Relations between Epistemology and Law

We want to end with the most general of doubts about much of Pardo's discussion – general enough to extend much further than the discussion of statistical evidence, or indeed of Safety and Sensitivity. For Pardo believes, perhaps like many others, that “[T]he process of legal proof is fundamentally, at root, an epistemic endeavor—an attempt to reach true conclusions based on evidence.” (5) We, on the other hand, think that epistemology doesn't matter legally. As theoretical disagreements go, there aren't many that are deeper⁴⁰.

⁴⁰ Still, we want to emphasize that many of our points against Pardo earlier in the text are independent of this wider disagreement.

We agree, of course, that *truth* matters legally. Truth, however, is not an epistemic concept at all: Whether a belief or a judgment is true does not depend (certainly not generally) about the available evidence for it, or the epistemic procedures we can use to find out whether or not it's true. So truth is beside the point. What we reject is the intrinsic legal significance of epistemic statuses like epistemic justification or knowledge – giving up any good the law can supply (say, in terms of incentives or justice) just in order to secure such a status for our judgments would be objectionably fetishistic⁴¹.

Think, for instance, about the instructions given to juries, or the presumption of innocence⁴². Now, these two give rise to questions and puzzles. And it may be thought that a way of making progress on them is by first making progress on the more general epistemological issues in the vicinity. Perhaps, for instance, the way to understand (or criticize) the epistemic-sounding instruction to juries or the presumption of innocence is to employ an epistemologically better-informed understanding of such concepts as *presumption* or *evidence* or *reasonable doubt*. But this would be a mistake. There are many contenders for considerations that normatively govern issues of legal evidence and procedure, but the thought that anything in the purely philosophical discussions of epistemology forces any legal answers upon us seems confused. Can you imagine, for instance, someone reasoning in the following way: “Sure, understanding the presumption of innocence in this specific way rather than that one will lead to less innocents in jail, or to better life-saving crime deterrence, but on the other hand, philosophical accounts of *presumption* pull in

⁴¹ Pardo (14, footnote 63) is right, of course, that even if Sensitivity doesn't matter to the law, it doesn't *follow* that no other epistemological property is. Though this doesn't follow (and we never suggested that it did), we do believe – on independent grounds – the more general claim as well.

In the same footnote, Pardo accuses us of failing to appreciate knowledge's factivity – so that knowledge entails truth. He thinks that this failure on our part undermines the thought experiment we use to support the claim that no epistemic credentials justify any price in any other legally relevant value. But he is wrong here. First, as noted in the text, we do (of course!) acknowledge a role for truth. So what we say about knowledge is only about whatever else is needed for knowledge, on top of truth – this other stuff is not intrinsically legally significant. Second, and most troublingly, Pardo gets close to misquoting us. He says (of one of the worlds we ask the reader to compare) “World A in which jurors “only convict when they know,” ...” noting that if they know that the accused is guilty, it follows that he is. Our text reads, however: “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...” (Enoch et al, supra note 5, at 212, italics added). We don't know why Pardo dropped the bracketed comment, but this is misleading editing on his part. Obviously, knowledge is factive (so is Safety, by the way, which may spell further trouble for Pardo). Just as obviously, though, thinking-that-one-knows is not factive (and this, of course, is why we introduced the bracketed comment to begin with).

⁴² We draw here on David Enoch *Comment on Yaffe's Attempts*, THE JERUSALEM REVIEW OF LEGAL STUDIES 6, 20-35 (2012, at 32-3).

the opposite direction, so hey, it's a tradeoff."? After all, if need be, we can distinguish between, say, presumption_e and presumption_l, let the philosophers argue about the former, and let the practical considerations that *are* legally relevant govern the latter. What would be lost?⁴³

It will remain, of course, a possibility that the epistemological and the legally relevant notions – in the example here, presumption_e and presumption_l – are after all related. But relying on this possibility requires substantive support, support that is not supplied merely by using a notation that hides the subscripts. And supplying such substantive support will require a discussion of whatever normative considerations govern the relevant parts of legal procedure and evidence law. We predict, of course – based on the fetishism point and the thought experiment from the previous paragraph – that no such substantive support can be supplied. But even if we're wrong about it, surely we're right that such support is needed. The thought that epistemology matters intrinsically to the law is, we think, false. But at the very least, it is not to be taken as a starting point.

⁴³ Here's David Lewis, taking for granted the point we argue for in the text here (and elsewhere): "What if some far-fetched possibility is called to our attention not by a sceptical philosopher, but by counsel for the defence? We of the jury may wish to ignore it, and wish it had not been mentioned. If we ignored it now, we would bend the rules of cooperative conversation; but we may have good reason to do exactly that. (*After all, what matters most to us as jurors is not whether we can truly be said to know; what really matters is what we should believe to what degree, and whether or not we should vote to convict.*)" (Our italics).

David Lewis, *Elusive Knowledge*, AUSTRALASIAN JOURNAL OF PHILOSOPHY 74, 549-567 (1996), at 560.