

‘More Likely Than Not’ Knowledge First and the Role of Bare Statistical Evidence in Courts of Law*

Michael Blome-Tillmann
University of Cambridge / McGill University

0. Abstract

The paper takes a closer look at the role of knowledge and evidence in legal theory. In particular, the paper examines a puzzle arising from the evidential standard *Preponderance of the Evidence* and its application in civil procedure. Legal scholars have argued since at least the 1940s that the rule of the Preponderance of the Evidence gives rise to a puzzle concerning the role of statistical evidence in judicial proceedings, sometimes referred to as the *Problem of Bare Statistical Evidence*. While this puzzle has led to the development of a multitude of accounts and approaches in the legal literature, I argue here that the problem can be resolved fairly straightforwardly within a knowledge-first framework.

1. The Puzzle of Bare Statistical Evidence

Proceedings in courts of law are governed by rules of evidence. Criminal proceedings in the USA and the UK, for instance, are governed the by rule of *Beyond Reasonable Doubt*, while the evidentially far less demanding rule of *Probable Cause* is applied to order searches, arrests, or to issue indictments in the mentioned jurisdictions. The evidential standard that I am interested in in this paper, however, governs civil procedure or tort suits—namely, the standard of the *Preponderance of the Evidence*.¹ This standard is usually paraphrased in legal textbooks as *more likely than not* or as *greater than 50% chance*. But we can be more precise. Since the mentioned paraphrases are meant to govern *evidential probabilities*, we may construe them as constraints on the probability of a proposition p conditional on the evidence available to and admissible in court. Let us therefore define the rule as follows:

Preponderance of the Evidence (PE):

p meets the standard of proof of PE iff $P(p|e) > .5$.

On this understanding, a proposition p has been proven to the standard of the preponderance of the evidence just in case the probability of p given the admissible evidence e is greater than .5. Thus, in order to find a defendant liable in a civil tort suit, the probability that the defendant is at fault must be greater than .5 given the evidence.

Consider an example for illustration:²

Red Cab Case – Version A:

Mr Smith is driving home late one night when his car is sideswiped and run off the road by a taxi. In the crash, Mr Smith suffers two broken legs. He decides to sue Red Cab Company for damages. His evidence is as follows: there are only two taxi compa-

* I am indebted to Brian Ball, Philip Ebert, Peter Graham, Bruno Guindon, Andrew Higgins, Alicia Hinarejos, Iwao Hirose, Jesper Kallestrup, Aidan McGlynn, Duncan Pritchard, Eric Schwitzgebel, Paulina Sliwa, Levi Spectre, Sarah Stroud, and Han-Ru Zhou. Special thanks go to Arif Ahmed, Matthew Kramer, and Martin Smith for extensive comments on an earlier version of this paper and to three anonymous referees for OUP. This work was supported by Marie Curie Actions (PIIF-GA-2012-328969 ‘Epistemic Vocabulary’).

¹ See (Ingram 2015) for an overview of the mentioned standards.

² See (Tribe 1971) for the original version of the example. The present formulation of the case is largely borrowed from (Thomson 1986: 199-200).

nies in town, Red Cab Company (all of whose taxis are red) and Blue Cab Company (all of whose taxis are blue). An independent bystander who witnessed the accident testifies that the taxi in question was red.

Given standard assumptions about the reliability of eyewitness testimony, the probability that the taxi was red, and thus that Red Cab Company is at fault, given her testimony, is well above the threshold of .5—namely, at roughly .7.³ Conditional on the evidence presented and admitted in court, it is therefore *more likely than not* that the car that caused Mr Smith’s injuries belonged to Red Cab Company. And since the standard of proof is met in the above fictitious case, Red Cab Company would, in the example, be found liable to pay damages to Mr Smith. The example illustrates a clear and uncontroversial application of the rule of PE.

However, consider a slight variation of the above example:

Red Cab Case – Version B:

Mr Smith is driving home late one night when his car is sideswiped and run off the road by a taxi. In the crash Mr Smith suffers two broken legs. He decides to sue Red Cab Company for damages. His evidence is as follows: there are only two taxi companies in town, Red Cab Company (all of whose taxis are red) and Blue Cab Company (all of whose taxis are blue), and of the taxis in town on the night of the accident 70% were operated by Red Cab Company and 30% by Blue Cab Company. This is the only evidence Mr Smith can produce against Red Cab Company, for even though he could see that it was a taxi that caused the accident, he couldn’t see its colour. No further eyewitnesses have come forward.

On evaluating the evidence in Version B of the red cab case it is again obvious that the evidential probability that Red Cab Company is at fault is well above the threshold of .5. Given the evidence—that is, given the fact that Red Cab Company’s market share on the night of the accident was 70%—the probability that the taxi that caused the accident was one of Red Cab Company’s is .7. However, despite the fact that PE—the standard of proof for civil cases—has been met in Version B of our example, courts routinely find, with overwhelming consistency, for the defendant in such cases. Moreover, I take it that such decision making accords rather well with our intuitions about the case, for most people are hesitant to impose liability on the basis of the market share evidence in Version B, citing the intuition that the evidence presented in court is insufficient for a verdict for the plaintiff.⁴

We thus face a puzzle. According to our intuitions, the court should find Red Cab Company liable in Version A of the Red Cab Case, but it shouldn’t do so in Version B—despite the fact that the standard of the Preponderance of the Evidence is clearly met in either case. What is more, the evidential probability that the defendant is at fault is identical in the two cases. Why, then, and on what basis, are civil courts willing to violate PE so blatantly?

The problem generalizes, as the following example, inspired by Jonathan Cohen’s (1977: 74-81) *Paradox of the Gatecrasher*, illustrates.

³ I assume a prior probability that the defendant is at fault of .5, which results—according to (Fields 2013: 1799)—in a posterior probability of even .77. As Fields (*ibid.*) points out, the probability that the defendant is at fault given positive identification falls below .5 only if the prior probability that the defendant is at fault is below roughly .3. See (Schauer 2003: 317, fn. 15) for further references on the reliability of eyewitness testimony.

⁴ See (Wells 1992) for empirical evidence to this effect.

The Gatecrasher – Version A:

The organizers of the local rodeo decide to sue John for gatecrashing their Sunday afternoon event. Their evidence is as follows: John attended the Sunday afternoon event—he was seen and photographed on the main ranks. No tickets were issued, so John cannot be expected to prove that he bought a ticket with a ticket stub. However, a local police officer observed John climbing the fence and taking a seat. The officer is willing to testify in court.

In this case, just as in the previous A-version of the Red Cab Case, courts will, without much hesitation, find for the plaintiff. The police officer's eyewitness testimony establishes, in the absence of countervailing evidence, with a high-enough probability that John gatecrashed. To illustrate this further, let us again assume that the police officer is, given the evidence presented in court, ordinarily reliable. Then, the probability, given the evidence, that John gatecrashed is .7—well above the threshold required by PE.⁵

However, there is again a very similar case in which courts would find differently. Consider what I shall call *Version B* of the gatecrasher example:

The Gatecrasher – Version B:

The organizers of the local rodeo decide to sue John for gatecrashing their Sunday afternoon event. Their evidence is as follows: John attended the Sunday afternoon event—he was seen and photographed on the main ranks during the event. No tickets were issued, so John cannot be expected to prove that he bought a ticket with a ticket stub. However, while 1,000 people were counted in the seats, only 300 paid for admission.

In this revised version of the case, courts will clearly find for the defendant. And again, I take it that this accords well with our intuitions about fairness and justice. It doesn't seem right to find John, who was randomly picked out in the arena, liable just because 70% of attendees at the rodeo gatecrashed. If such a case was allowed to succeed, the organizers of the rodeo could, after all, in principle win similar cases for every person in attendance at the rodeo, including the 300 people that paid the entrance fee.⁶

We are again facing a puzzle. What justifies courts' finding differently in the A and the B-version of the Gatecrasher, given that the evidential probability that the defendant is at fault is in either case well above the threshold required by the rule of the Preponderance of the Evidence? Given that all evidence is probabilistic and that the evidential probabilities are identical with respect to the A and B-versions of our cases, courts should find for the plaintiff in the A-versions if and only if they do so in the B-versions. But the imposition of liability on the basis of bare statistical evidence in the B-versions is intuitively unwarranted and unjust. But, then, what is, in the light of the rule of the Preponderance of the Evidence, the decisive difference between the two different types of cases?

2. Individual vs. Bare Statistical Evidence

Cases such as the Red Cab/Blue Cab example and the Gatecrasher have puzzled lawyers, judges, and jurists since at least the 1940s. A number of explanations have been

⁵ (Fields 2013: 1799).

⁶ To soften the blow one might restrict compensation to actual losses and share liability amongst all attendees at the rodeo. While slightly less unfair to those individuals who paid the entrance fee, I take it that such an approach would still place unreasonable demands on fee-paying attendees. Thanks to an anonymous referee for OUP here.

proposed to resolve the puzzle.⁷ To begin with, let us distinguish two different broad approaches to the puzzle, what I shall call the *conservative* and the *revisionist* approaches. According to the revisionists, our willingness to violate PE in the B-versions of the above cases is due to a flawed underestimation of the evidential standing of statistical evidence. On this view we are mistaken in “excluding more-reliable statistical evidence and admitting less-reliable personal testimony, [a practise that] is explained by the widespread but empirically unsupported faith in eyewitness identification” (Schauer 2003: 94). I shall not discuss such revisionist approaches in this paper. Rather, I take the intuition that courts would be at fault in finding the defendant liable in the above B-cases to be far too strong a datum to be simply ignored in favour of revisionism. Instead, I shall here explore a conservative approach to the puzzle that aims to explain why courts are right in finding the way they actually do.

A large number of jurists have argued along conservative lines that the puzzle of bare statistical evidence can be resolved by distinguishing between two fundamentally different types of evidence. We can find such distinctions, drawn with more or less rigour, in a large number of court judgments, but the distinction is also widely discussed in the academic literature on evidence law.⁸ The distinction invoked is often referred to as that between *individual* and *bare statistical* evidence.⁹ JJ Thomson (1986), for instance, observes that

what is at work in the friends of individualized evidence is precisely the feeling that just imposition of liability requires that [a] stronger requirement be met. They believe [...] that “mathematical chances” or “quantitative probability” is not by itself enough; on my view of them, that is because they feel, rightly, that if a jury declares a defendant guilty on the ground of nonindividualized evidence alone, then it is just luck for the jury if what it declares true is true—and they feel, not without reason, that it is unjust to impose liability where that is the case. (Thomson 1986: 214)

As Thomson points out, bare statistical evidence doesn’t seem to eliminate luck. Intuitively, if one’s judgement is based on bare statistical evidence, then one’s judgement is, if correct, correct as matter of luck. Judgements that are based on individualized evidence, such as eyewitness testimony, however, are intuitively correct not merely as a matter of luck. Individual evidence, in short, seems to eliminate an element of luck that bare statistical evidence fails to rule out.

The issue can be illustrated by an analogy to lottery examples. Even though the evidential probability that one’s lottery ticket is a loser may be exceedingly high, one’s belief that it is a loser is, if true, true as a matter of luck: one’s belief that one’s ticket is a loser could, after all, very easily have been false (the balls deciding the winning numbers could very easily have fallen in a different order, making one’s ticket a winner). Thus, for all one knows, one’s ticket might win the jackpot, and even though the probabilities speak strongly against such a positive result, one’s belief that one won’t win is, in an intuitive sense, true as a matter of luck. What the lottery example and the B-versions of the above examples have in common is thus that a high evidential probability alone cannot eliminate luck. The elimination of the relevant type of luck, however, seems to be what is necessary for a proposition to meet the evidential standards that are in fact applied in civil tort suits.

⁷ See, for instance, *Smith v Rapid Transit, Inc.*

⁸ See (Thomson 1986: 200) for references.

⁹ Some theorists and judges use the terms ‘direct evidence’ and ‘general evidence’ (see *Smith v Rapid Transit, Inc.*).

Before taking a closer look at how to illuminate the distinction between individual and bare statistical evidence in the following section, note that the conservative strategy under consideration leads to a reformulation of (PE). Here is (PEI), where ‘ e_i ’ refers to the court’s individual evidence:

Preponderance of the Individual Evidence (PEI):
 p meets the standard of proof in civil cases iff $P(p|e_i) > .5$.

While this strategy—and the distinction underlying it—is intuitively appealing, it is in need of further clarification. We do not yet have a very clear understanding of the difference between individual and bare statistical evidence. There have been several attempts to elucidate the distinction in the recent literature. Cohen (1977), for instance, has argued that the notion of probability at issue is not that of ordinary evidential probability but rather that of *inductive probability*. JJ Thomson (1986) has offered an account of individual evidence in causal terms, and, more recently, Enoch et al. (2012) have argued that individual evidence can be accounted for in terms of modal sensitivity. Discussion of these approaches, however, is beyond the scope of this paper.¹⁰ Instead, I shall here propose a novel approach to the distinction between individual and bare statistical evidence that aims to resolve the puzzle of bare statistical evidence by resorting to the conceptual framework of a knowledge-first epistemology.

3. The Knowledge Account

A recent revolution in epistemology—the knowledge-first movement—favours knowledge as the basic notion of epistemological theory. Knowledge is, on this view, taken as a primitive, which is then used to elucidate and explain a number of other philosophically interesting concepts and phenomena. Moreover, a number of knowledge-firsters also take the view that knowledge plays an important normative role. On such views knowledge is, for instance, the norm of practical reasoning, assertion, or belief. Consider, for illustration, the knowledge rule of assertion as proposed by Timothy Williamson (2000: 243):

(TW) It is permissible for x to assert p only if x knows p .

While (TW) is the principle most widely associated with the label ‘Knowledge Rule of Assertion’, other authors—notably Hawthorne (2004: 23, fn. 58) and DeRose (2002: 167, 187)—have been tempted by a logically stronger, biconditional version of Williamson’s knowledge rule:¹¹

(A) It is permissible for x to assert p iff x knows p .

Along similar lines, Hawthorne (2004: 30) and Hawthorne and Stanley (2008) have defended the view that knowledge is the norm of practical reasoning. Here is the *Knowledge Rule of Practical Reasoning*:

(PR) It is permissible for x to use p as a premise in her practical reasoning iff x knows p .

¹⁰ But see (Blome-Tillmann forthcoming) for criticism of the mentioned views.

¹¹ Cp. (Brown 2010: 249-250).

Drawing inspiration from Williamson (2005: 108) and McGrath (2010: 396), we may add the *Knowledge Norm of Belief* to the mix:

- (B) It is permissible for x to outright believe p iff x knows p .¹²

Several authors have objected to one or another of these principles on different grounds.¹³ However, the knowledge-first approach is growing in popularity and promises hope for a slow but steady paradigm shift in epistemology.

What is the relevance of the above views about the normative role of knowledge for the puzzle of bare statistical evidence? Assuming that we ought to act on p only if we know p , the ideal and strongest epistemic state a court could be in, from a normative point of view, is one in which the court knows whether the defendant is at fault. It is, for instance, only if the court knows that the defendant is at fault that the court should assert, outright believe, or use as a premise in its judicial decision making the proposition that the defendant is at fault. However, we certainly cannot require courts to reach the epistemically normative ideal state of knowledge in the sense that a judgement for the plaintiff can only be delivered, if the court knows that the defendant is at fault. Courts need to be able to make decisions in the absence of the epistemically normative ideal—that is, in the absence of knowledge that the defendant is at fault.

What is needed is a measure of gradual approximation to the normative ideal. One way to measure such approximation is in terms of *evidential probabilities that one knows* that the defendant is at fault. Knowledge then remains the normatively ideal state, but epistemic success in courts of law can be measured and understood in terms of something less than knowledge—namely, the evidential probability that knowledge has been achieved. On the view proposed here, courts are thus to measure the evidential probability that they are in the epistemically ideal state with respect to the proposition that the defendant is at fault rather than simply the evidential probability that the defendant is at fault.¹⁴ What matters, in civil cases, is thus $P(Kp|e)$ rather than $P(p|e)$. The view at issue is a natural extension of the knowledge first approach in epistemology to legal theory and evidence law.

Consider a second argument for the knowledge approach to evidence in judicial proceedings. In courts of law, we want to get things right, and—as we have seen above—we do not just want to get things right by luck. In other words, if a court takes p to be established according the relevant rules of evidence, then, ideally, we want it to be the case that p is, if true, not just true as a matter of luck. Rather, we want to be sure, to varying degrees, that we are not lucky in getting things right, if we are getting them right. What is needed is not just evidence that we got it right, but also evidence that we are not merely lucky in getting it right. Next, note that it is a familiar fact that if one knows p , then one is not just lucky in getting it right with respect to p —knowledge precludes epistemic luck (as is familiar from the lottery cases briefly men-

¹² As Williamson (2005: 108) puts it, ‘if one knows [p], then one can hardly be wrong to believe [p]; conversely, given that one does not know [p], it arguably is wrong to believe [p].’ McGrath (2010: 396) only assumes the sufficiency of knowledge for permissible belief.

¹³ See, for instance, (Brown 2008, 2010, Goldman 2008, Kvanvig 2007, Lackey 2007, Neta 2009). For an alternative view to the knowledge norm of belief, see (Wedgwood 2002, 2013).

¹⁴ Note also that the view proposed is not at odds with the above norms and, in particular, not with (PR). To see this, note that the idea that it is permissible to rely on p in our practical reasoning iff we know p remains satisfied. The court should only rely on a proposition p if it knows p —which means, in the particular cases envisaged, that the court should rely on propositions about evidential probabilities iff it knows those propositions. These constraints are satisfied in proceedings that are not in violation of due process.

tioned above).¹⁵ So what is needed in court is evidence that one knows that the defendant is at fault and not just evidence that the defendant is at fault.¹⁶

Let us be more precise about the view proposed here. There are at least two different ways to implement the main idea underlying the above considerations. On the first approach, we replace the traditional rule of the Preponderance of the Evidence with what I shall call the *Knowledge Rule of the Preponderance of the Evidence* (KPE):

Knowledge Rule of the Preponderance of the Individual Evidence (KPE):
 p meets the standard of proof in civil cases iff $P(Kp|e) > .5$.

According to (KPE), the standard of proof is met in a civil case iff the evidential probability that the court knows that the defendant is at fault is greater than .5. This view is revisionary to some degree, as it replaces (PE) with (KPE), but it is—as we shall see below—in harmony with our intuitions about the just imposition of liability in courts. A second, slightly different approach leaves the original (PE) in place but formulates a novel constraint on what type of evidence is admissible in court. Consider the following principle:

Evidence Constraint (EC):
 e is admissible in court as evidence for p only if it raises the probability that someone in possession of e knows p .

(EC) allows us to achieve the same results as (KPE) by disqualifying what we have previously called ‘bare statistical evidence’ from civil proceedings—without amending the familiar rule of the Preponderance of the Evidence (PE). As far as I can see, both of these approaches are promising.

Next, consider how the above approaches resolve the puzzle of bare statistical evidence. Consider first the Gatecrasher cases. In the A-version of the example, the probability that the court knows that John is at fault is reasonably high given the court’s evidence. Given that the average reliability of witnesses in court, the probability that the witness in this case has correctly identified a gatecrasher is, given the court’s evidence, roughly .7.¹⁷ Since there is no countervailing evidence to the effect that the court has been led astray, the probability that the court knows that John is at fault is high: something unusual would have to have happened for the court not to know that the defendant is at fault.¹⁸ Of course, the court fails to know that they know

¹⁵ See also (Pritchard 2005) for the connection between knowledge and luck.

¹⁶ Cp. *Day v. Boston & Maine Railroad*, 96 Me. at 207, 52 A. at 774, quoted in Thomson, “Liability and Individualized Evidence,” p. 234, n. 13: “Quantitative probability, however, is only the greater chance. It is not proof, nor even probative evidence, of the proposition to be proved. That in one throw of dice there is a quantitative probability, or greater chance, that a less number of spots than six will fall is no evidence whatever that in a given throw such was the actual result. Without something more, the actual result of the throw *would still be utterly unknown*. The slightest real evidence that sixes did in fact fall uppermost would outweigh all the probability otherwise”. As Thomson (1986: 206, n. 14) points out, the passage is somewhat obscure, as it suggests that statistical evidence doesn’t probabilify. However, it also nicely brings to light the legal relevance of the intuition that knowledge is incompatible with epistemic luck.

¹⁷ See (Fields 2013: 1799).

¹⁸ The probability that nothing unusual has happened is, in the above cases, measured by the reliability figures of witness testimony in court. Thus, the evidential probability that nothing unusual has happened and that the defendant is in fact at fault, as testified by the witness, is roughly .7. Consequently, the probability that the court knows is roughly .7.

that the defendant is at fault, but the evidential probability that the court knows is, given the court's evidence, fairly high. Therefore, the court ought to find for the plaintiff in the A-version of the Gatecrasher.

In the B-version of the Gatecrasher, however, things are rather different. In this case, the court's belief that John gatecrashed is, if true, true as a matter of luck. For all the court knows, John could have easily been one of the 300 attendees at the rodeo who paid the entrance fee and there is no evidence whatsoever to eliminate that hypothesis. Thus, the court knows (or is at least in a position to know) that it does not know that John is at fault. In other words, given the evidence available to the court, the probability that the court knows that John is at fault is zero—hence the strict and categorical difference to our intuitions in the A-version of the Gatecrasher. Imposing liability in the B-version of the case would be gravely inappropriate, and the strength of that intuition is accounted for by the fact that the probability that the court knows that the defendant is at fault is zero in the B-version of the Gatecrasher.

Next, consider the Red Cab cases. In the A-version of the example, the probability that the court knows that Red Cab Company is at fault is reasonably high given the evidence. Given the average reliability of the type of witness at issue, the probability that the witness has correctly identified the colour of the taxi is, given the court's evidence, .7. Since there are no reasons to assume that the court has been led astray, the probability that the court knows that the taxi was red is high: something unusual would have had to happen for the court not to have such knowledge. The probability that nothing unusual has happened, however, is measured by the reliability figures of witnesses in court. Thus, the evidential probability that the court knows that the defendant is at fault is roughly .7 and thus above the threshold required by the knowledge rule of the preponderance of the evidence.¹⁹

In the B-version of the case, however, things are again rather different. To see this, note again that in the B-version of the Red Cab case, the court knows (or is in a position to know) that it does not know that the defendant is at fault. Thus, the probability that the court knows that the defendant is at fault given the court's evidence is zero—analogously to the B-version of the Gatecrasher example. Again, the knowledge account explains rather elegantly our intuition that it would be grossly unjust for the court to find for the plaintiff in the B-version of the Red Cab case.

In summary, the knowledge approach to individual evidence does not only cohere rather well with a knowledge-first approach in epistemology, it also explains our reluctance to impose liability in cases involving bare statistical evidence. Individual evidence, after all, is crucially not more potent in virtue of being more highly probabifying than bare statistical evidence. Rather, individual evidence is perceived to be more potent in courts of law because it has the power to raise the probability that one knows. Evidential probability, in the broad sense, that the defendant is at fault is therefore not really what matters in civil cases. Rather, it is the evidential probability that the court knows the defendant to be at fault. It is the latter type of probability that tracks our intuitions concerning the just imposition of liability.

4. Some Objections

4.1. Knowledge and Safety

It is worthwhile illustrating the above considerations further by means of the Safety principle. Why is the evidential probability that the court knows that the defendant is

¹⁹ See (Fields 2013: 1799).

at fault zero in the B-versions of the above cases? To illustrate this further, consider Timothy Williamson's (2000: 147) preferred formulation of Safety:

(WS) If one knows, one could not easily have been wrong in a similar case.²⁰

(WS) is independently motivated, but I shall not go into the details of the literature on safety in this paper.²¹ Instead, note that (WS) allows us to explain why $P(Kp|e) = 0$ in the B-versions of the above cases. To see this, note that in those cases the court's belief that the defendant is at fault, if it forms such a belief, is clearly not safe. For the court could, in the B-versions of our examples, easily have been wrong in a similar case—namely, in all those similar cases in which the taxi at issue was blue (in the Red Cab example) or in which John paid the entrance fee to the rodeo (in the Gatecrasher example). Those possibilities are very similar to the ones described in the relevant B-versions, but it is important to point out that the court is mistaken in believing, in those rather similar cases, that the defendant is at fault. Thus, (WS) explains elegantly why the court does not know that the defendant is at fault in the B-versions.

What is more, (WS) also allows us to account for the fact that the court knows (or is in a position to know) that they do not know that the defendant is at fault. That is so because it is rather obvious that the court could very easily have been wrong in a similar case. It is obvious, in other words, that the court's belief that the defendant is at fault, if it has formed such a belief, is not safe. But if the court knows that they do not know that the defendant is at fault because their corresponding belief is not safe, then the proposition that they do not know is part of the court's evidence, and $P(Kp|e)$ must, therefore, equal zero.²² In summary, (WS) explains why courts cannot acquire knowledge on the basis of bare statistical evidence, and thus helps motivating the view that $P(Kp|e) = 0$. Thus, (WS) provides us, in conjunction with (KPE) or (EC), with a good explanation of why we cannot use bare statistical evidence in courts of law.²³

It might be objected at this point that it would be methodologically simpler to resolve the puzzle of bare statistical evidence by appeal to safety rather than knowledge.²⁴ Shouldn't we measure the evidential probability that the court's belief is safe rather than the evidential probability that it knows when contemplating the just imposition of liability? Nothing much hangs on the answer to this question. However, I take it that, for independent (and largely methodological) reasons, our only grasp of the notion of safety is via the notion of knowledge: safety, or the notion of similarity between cases employed in (WS), cannot be defined in terms other than knowledge.²⁵ It is, therefore, knowledge—not safety—that comes explanatorily first. On such an approach, what is really doing the explanatory work is the notion of knowledge—even if we were to reformulate the rule of the Preponderance of the Evidence in terms of Safety. Thus, on the knowledge first view, the methodologically simpler approach is to measure the probability that the court knows rather than the probability that the

²⁰ In his (2009: 23) Williamson illustrates this view further by arguing that “knowing p requires safety from the falsity of p and of its epistemic counterparts.”

²¹ But see (Pritchard 2005, 2009, Williamson 2000, 2009).

²² I here assume that if one knows p , then p is part of one's evidence. This principle is entailed by, but not equivalent, to Timothy Williamson's (2000) thesis that $E=K$. See also (Blome-Tillmann 2015) for discussion.

²³ In the A-versions of our cases, on the other hand, the evidence can be interpreted as raising the probability that the court's relevant beliefs are safe, as outlined for knowledge in the previous section.

²⁴ See (Pritchard ms) for an account of legal risk in terms of safety.

²⁵ Cp. (Blome-Tillmann 2014, ch. 5).

court's beliefs are safe. It is knowledge—not safety—that is the ultimate anti-luck condition in epistemology.²⁶

4.2. *The Role of Knowledge in Legal Theorizing*

It might be argued that knowledge is a confusing or somewhat slippery notion and that it therefore ought not to play a role in courts of law or our theories of evidence law.²⁷ But it is important to note that the defenders of such a view still owe us an account of evidence—a *prima facie* equally confusing and slippery notion. They must, in other words, provide us with an answer to the question: under what conditions is a proposition *p* part of the court's evidence? One simple and explanatorily powerful response to this challenge derives from Timothy Williamson's view that $E=K$ —that is, from the view that all and only those propositions that one knows are part of one's evidence. On such a view, a proposition *p* is part of the court's evidence iff the court knows *p*. Now, if we accept that $E=K$, then we are stuck with the notion of knowledge in our theorizing about evidence law independently of our acceptance of (KPE) or (EC). If we reject $E=K$, on the other hand, then we must offer an alternative account of evidence—an account that is not only independently well motivated (as $E=K$ is) but also workable within judicial contexts. While it surely is not impossible to develop such a theory of evidence, rejecting a knowledge-based account of evidence in legal contexts on the worry that the notion of knowledge may appear somewhat slippery will seem to many to amount to throwing the baby out with the bathwater.

4.3. *Negligence, Racial Discrimination, and Statistical Evidence as Knowledge-Generating*

Sometimes, but only sometimes, we can know on the basis of bare statistical evidence. Some seemingly bare statistical evidence leads to knowledge.²⁸ Thomson (1986: 205) points out, for instance, that in racial discrimination suits, evidence about the racial composition of both an organization's work force and the local population may be admissible and lead to a verdict for the plaintiff. It is noteworthy, however, that in such cases we also intuit that—if the statistical evidence is strong enough—we *know* that the organization discriminated racially (or at least that the probability that we know given our evidence is significantly higher than .5).²⁹ Thus, as long as our intuitions about whether we know correlate with our intuitions about whether statistical evidence is sufficient for the just imposition of liability, the mere fact that statistical evidence is sometimes sufficient for the just imposition of liability does not threaten the view defended here.³⁰

Further cases in which bare statistical evidence appears sufficient for the just imposition of liability comprise cases of so-called *market share liability*.³¹ In certain pharmaceutical or tobacco litigation, liability is assigned according to the market share a particular company had—even if the company's market share is below the threshold of 50% and thus reduces the evidential probability that the company is re-

²⁶ On the view proposed here, safety is eventually nothing but a heuristic device that illustrates a property of knowledge—a property that cannot be understood or defined independently of knowledge.

²⁷ Enoch et al. (2012) seem to be attracted by such a view.

²⁸ See (Blome-Tillmann 2014, ch. 5) for discussion.

²⁹ Thomson's view on the matter is that the statistical evidence is in this case causally explained by the fact that there is actual racial discrimination.

³⁰ Further cases of the abovementioned type are, presumably, some cases involving DNA evidence. I shall leave the treatment of issues arising from the admissibility of DNA evidence in courts of law for another occasion.

³¹ See *Sindell v. Abbott Laboratories* (1980) 26 C3d 588.

sponsible for the plaintiff's impairments to below .5. However, it is important to emphasize that in successful cases of market share liability the negligence of the defendant companies has usually been established independently. Thus, liability is assigned according to market share only once the defendant has been independently shown, by the standard of the preponderance of the evidence, to have acted negligently. Consequently, cases of market share liability provide harmless exceptions to the principles of (KPE) and (EC) above. The relevant examples can easily be accounted for by restricting (KPE) and (EC) to cases in which negligence has not been established independently.

4.4. Causation as an Epistemic Proxy

In an important paper on the puzzle of bare statistical evidence JJ Thomson (1986) argues that individual evidence ought to be understood in causal terms. According to Thomson,

(JJT) e is individual evidence in support of p iff the fact that p caused e .

I have argued elsewhere against Thomson's view, but its usefulness must be emphasized here.³² To see what I have in mind note the following connection between knowledge and causation:

(KC) If x knows p , then x 's belief that p was caused by the fact that p .

(KC) is true with respect to most, if not all, cases of empirical knowledge. If (KC) is true, however, then we can use causation as an indicator of knowledge. In other words, from a practical point of view, courts may aim to determine whether the above causal constraint is met with sufficient evidential probability, instead of aiming to settle the potentially harder question of whether (KPE) is satisfied. While the theoretical explanation of the puzzle must rest on the knowledge account outlined above, the practical endeavour of decision making in courts of law can safely be guided by the causal account outlined by Thomson. While causation by itself is not sufficient to eliminate epistemic luck, it is, in ordinary cases (including the cases discussed in Section 1 of this article), a reliable indicator of its elimination and thus of knowledge.

5. Conclusion

Knowledge matters. We value knowledge that p higher than the mere high evidential probability that p . Moreover, in many cases of legal relevance—such as the ones discussed in this paper—we value a high evidential probability that we know that p significantly more highly than the mere high evidential probability that p . The actual behaviour of courts and their dismissal of bare statistical evidence is a manifestation of this fact: our intuitions about the just imposition of liability in civil suits can be accounted for in a rather straightforward manner within a knowledge first approach to legal standards of proof. It is, therefore, knowledge, along with our appreciation of the likelihood that we have it, that guides our actions in courts.

³² One obvious problem with (JJT) is that causation is factive. Thus, given (JJT) a judgement meets the standard of proof of the preponderance of the evidence only if the defendant is in fact at fault. This seems an unintuitive and problematic consequence of (JJT). For a more detailed discussion of Thomson's view see (Blome-Tillmann forthcoming).

6. References

- Blome-Tillmann, M. (2014). Knowledge and Presuppositions. Oxford, Oxford University Press.
- (2015). "Solving the Moorean Puzzle." Philosophical Studies **172**(2): 493-514.
- (forthcoming). "Sensitivity, Causality, and Statistical Evidence in Courts of Law." Thought.
- Brown, J. (2008). "The Knowledge Norm of Assertion." Philosophical Issues **18**: 89-103.
- (2010). "Knowledge and Assertion." Philosophy and Phenomenological Research **81**(3): 549-566.
- Cohen, L. J. (1977). The probable and the provable. Oxford, Clarendon Press.
- DeRose, K. (2002). "Assertion, Knowledge, and Context." The Philosophical Review **111**(2): 167-203.
- Enoch, D., L. Spectre and T. Fisher (2012). "Statistical Evidence, Sensitivity, and the Legal Value of Knowledge." Philosophy & Public Affairs **40**(3): 197-224.
- Fields, K., L. (2013). "Toward a Bayesian Analysis of Recanted Eyewitness Identification Testimony." New York University Law Review **88**(1769): 1769-1801.
- Goldman, A. (2008). "Knowledge, Explanation, and Lotteries." Noûs **42**(3): 466-481.
- Hawthorne, J. (2004). Knowledge and Lotteries. Oxford, OUP.
- Hawthorne, J. and J. Stanley (2008). "Knowledge and Action." Journal of Philosophy **105**(10): 571-590.
- Ingram, J. L. (2015). Criminal Evidence. Waltham, MA, Anderson Publishing.
- Kvanvig, J. L. (2007). "Contextualism, Contrastivism, Relevant Alternatives, and Closure." Philosophical Studies **134**: 131-140.
- Lackey, J. (2007). "Norms of Assertion." Noûs **41**(4): 594-626.
- McGrath, M. (2010). "Contextualism and Intellectualism." Philosophical Perspectives **24**(1): 383-405.
- Neta, R. (2009). "Treating Something as a Reason for Action." Noûs **43**(4): 684-699.
- Pritchard, D. (2005). Epistemic Luck. Oxford, Clarendon.
- (2009). "Safety-Based Epistemology: Whither Now?" Journal of Philosophical Research **34**: 33-45.
- (ms). "Risk."
- Schauer, F. (2003). Profiles, Probabilities, and Stereotypes. Cambridge, Massachusetts, Harvard University Press.
- Thomson, J. J. (1986). "Liability and Individualized Evidence." Law and Contemporary Problems **49**(3): 199-219.
- Tribe, L. H. (1971). "Trial by mathematics: precision and ritual in the legal process." Harvard Law Review **94**(6): 1329-1393.
- Wedgwood, R. (2002). "The Aim Of Belief." Noûs **36**: 267-297.
- (2013). The Right Thing to Believe. The Aim of Belief. T. Chan, Oxford University Press: 123-139.
- Wells, G. L. (1992). "Naked Statistical Evidence of Liability: Is Subjective Probability Enough?" Journal of Personality and Social Psychology **62**(5): 739-752.
- Williamson, T. (2000). Knowledge and Its Limits. Oxford, OUP.
- (2005). Knowledge, Context and the Agent's Point of View. Contextualism in Philosophy: On Epistemology, Language and Truth. G. Preyer and G. Peter. Oxford, OUP: 91-114.
- (2009). "Probability and Danger." The Amherst Lecture in Philosophy **4**: 1-35.

