

LIABILITY AND INDIVIDUALIZED EVIDENCE

JUDITH JARVIS THOMSON*

I

Cases like *Smith v. Rapid Transit, Inc.*¹ present a problem to students of tort law. Here is a typical hypothetical case—I will call it *Smith v. Red Cab*—which presents the problem more cleanly than the actual case does. Mrs. Smith was driving home late one night. A taxi came towards her, weaving wildly from side to side across the road. She had to swerve to avoid it; her swerve took her into a parked car; in the crash, she suffered two broken legs. Mrs. Smith therefore sued Red Cab Company. Her evidence is as follows: she could see that it was a cab which caused her accident by weaving wildly across the road, and there are only two cab companies in town, Red Cab (all of whose cabs are red) and Green Cab (all of whose cabs are green), and of the cabs in town that night, six out of ten were operated by Red Cab. Why is that the only evidence she can produce against Red Cab? She says that although she could see that it was a cab which came at her, she could not see its color, and as it was late, there were no other witnesses to the accident—other than the driver himself, of course, but he has not come forward to confess.

If we believe Mrs. Smith's story, and are aware of no further facts that bear on the case, then we shall think it .6 probable that her accident was caused by a cab operated by Red Cab. I think it pays to spell this reasoning out; what follows is one way of doing so. If we believe Mrs. Smith's story, then we believe that a cab, indeed exactly one cab, caused the accident, so that there is such a thing as *the* cab which caused the accident; and we believe that it was a cab in town that night. Thus we believe:

- (1) The cab which caused the accident was a cab in town that night.

If we believe Mrs. Smith's story, we also believe:

- (2) six out of ten of the cabs in town that night were operated by Red Cab.

Relative to the facts reported by (1) and (2), it is true that:

- (3) The probability that the cab which caused the accident was operated by Red Cab is .6.

Copyright © 1986 by Judith Jarvis Thomson

* Professor of Philosophy, Massachusetts Institute of Technology. A great many people made helpful criticisms of earlier drafts of this essay; I am particularly indebted to Jonathan Bennett, David Gauthier, Gilbert Harman, Paul Horwich, and Mary C. Potter.

1. 317 Mass. 469, 58 N.E.2d 754 (1945). Mrs. Smith's evidence against the defendant bus company consisted of evidence that she was caused harm by a negligently-driven bus on Main Street, and that the defendant bus company has the sole franchise for operating a bus line on Main Street.

But those are the only facts such that we are both aware of them and aware of their bearing on the question who operated the cab which caused the accident. (Perhaps we are aware that the accident took place on, as it might be, a Tuesday. Even so, we are not aware of any reason to think that fact bears on the question whose cab caused the accident.) Other facts whose relevance is clear might come out later: for example, a Green Cab driver might later confess. But as things stand, we have no more reason (indeed we have less reason) to think that any facts which later come out would support the hypothesis that the cab which caused the accident was operated by Green Cab than we have to think they would support the hypothesis that the cab which caused the accident was operated by Red Cab. We are therefore entitled to conclude that (3) is true—in fact, rationality requires us to conclude that (3) is true, for .6 is the degree of belief that, situated as we are, we ought to have in the hypothesis that the cab which caused the accident was operated by Red Cab.²

Is it right that Mrs. Smith win her suit against Red Cab? The standard of proof in a tort suit is “more probable than not,”³ which is plausibly interpretable as requiring only that the plaintiff establish a greater than .5 probability that the defendant (wrongfully) caused the harm. But most people feel uncomfortable at the idea of imposing liability on Red Cab on such evidence as Mrs. Smith here presents. Why? That is the problem.

II

*People v. Collins*⁴ and its typical descendant hypotheticals raise an analogous problem for the student of criminal law, but less cleanly, so let us set them aside for the time being. Consider, instead, a hypothetical case which I shall call *People v. Tice*.⁵ Two people, Tice and Simonson, both hated Summers and wished him dead. Summers went hunting one day. Tice followed with a shotgun loaded with ninety-five pellets. Quite independently, Simonson also followed, but *he* had loaded his shotgun with only five pellets, that being all he had on hand. Both caught sight of Summers at the same time, and both shot all their pellets at him. Independently: I stress that there

2. Some commentators on earlier versions of this essay have suggested that perhaps we ought not be entirely confident of the truth of (3), given only the truth of (1) and (2), on the ground that there just might have been more green cabs than red cabs in the part of town in which, and at the time at which, the accident took place. This suggestion seems to me to be a mistake—but one which is not worth going into here. Anyone who is similarly moved may suppose we are provided, in addition to (1) and (2), with an admission by Red Cab that there is good reason to think (and no reason not to think) that cab-distribution at the time, in the part of town south of the tracks (where the accident took place), was the same as overall cab-distribution in town that night.

3. See generally W. PROSSER & W.P. KEATON, PROSSER & KEATON ON THE LAW OF TORTS § 38 (5th ed. 1984).

4. 68 Cal. 2d 319, 438 P.2d 33, 66 Cal. Rptr. 497 (1968). Collins and his wife had been convicted of second degree robbery on evidence that the robbers, like the defendants, were an interracial couple; the man black, with mustache and beard, who drove a yellow car, the woman white, with blond hair in a ponytail—the prosecution alleged on statistical grounds that there was only one chance in twelve million that a couple would possess all these features.

5. With apologies to *Summers v. Tice*, 33 Cal. 2d 80, 199 P.2d 1 (1948).

was no plot or plan. Only one pellet hit Summers, but that one was enough: it hit Summers in the head and caused his death. While it was possible to tell that the pellet which caused Summers' death came either from Tice's gun or from Simonson's gun, it was not possible to tell which. So what charges should be brought against Tice and Simonson? In the event, Simonson is charged with attempted murder, and in *People v. Tice*, Tice is charged with murder.

Well, why not? To win its case against Tice, the prosecution must show that it is beyond a reasonable doubt that the pellet which caused Summers' death was a pellet fired by Tice. But given the information in hand, that seems easy, for given the information in hand, we can say both:

(1') The pellet which caused Summers' death was a pellet fired at Summers

and

(2') Ninety-five out of the one hundred pellets fired at Summers were fired by Tice.

The facts these report are the only facts such that we are both aware of them and aware of their bearing on the question whose pellet caused Summers' death. We therefore may, indeed should, conclude that it is also true that:

(3') The probability that the pellet that caused Summers' death was fired by Tice is .95.

And isn't a proposition beyond a reasonable doubt if it is .95 probable?⁶

I hope you will feel at least as uncomfortable at the idea of convicting Tice of murder on such evidence as this that he caused the death as you feel at the idea of imposing liability on Red Cab in *Smith v. Red Cab*.

There are differences, of course. In *Smith v. Red Cab*, the information we have in hand gives no reason at all to think that both cab companies were at fault: it gives reason to think that one cab company was at fault, namely the one, whichever it was, that caused the accident. So if Mrs. Smith wins her suit, then Red Cab may be being held liable for her costs despite the fact that it not only did not cause her injury, but was entirely without fault.

By contrast, the information we have in hand in *People v. Tice* gives reason to think that both Tice and Simonson were at fault: it gives reason to think that both committed attempted murder. So if the people win their case against Tice, then while Tice may be being held liable for murder without having caused Summers' death, he was all the same gravely at fault, having at a minimum tried to bring that death about.

This difference brings in train yet another. If Tice did not cause Summers' death, then his failure to do so was—relative to the evidence we have in hand—just *luck*, good or bad luck according to the view you take of the matter. He did everything he could to cause the death, and if he did not cause it, well,

6. Rita James Simon invited a sample of judges to translate "beyond a reasonable doubt" into numerical probabilities. It is interesting to learn that almost a third of them translated it as a probability of one. Over two thirds, however, translated it as a probability of .95 or less. See Simon, *Judges' Translations of Burdens of Proof into Statements of Probability*, 13 TRIAL LAWYER'S GUIDE 103 (1969).

that was certainly no credit to him. By contrast, if Red Cab did not cause Mrs. Smith's injury, that was a credit to Red Cab, for—relative to the evidence we have in hand—if Red Cab did not cause the accident, it was not at fault at all.

We could have eliminated these differences by altering the details of *People v. Tice* so as to make the evidence suggest only that one of the actors (Tice or Simonson) was at fault, the evidence that it was Tice who was at fault issuing from nothing other than the evidence that it was a pellet from Tice's gun which caused the death—just as the evidence that it was Red Cab that was at fault issues from nothing other than the evidence that it was one of Red Cab's cabs which caused Mrs. Smith's accident. It is not easy to alter the details in that way without introducing a measure of weirdness. But I think we ought to feel that there is no need to do so; that is, I think we ought not be moved by the differences I pointed to between *Smith v. Red Cab* and *People v. Tice*.

No doubt it was just luck for Tice if he did not cause Summers's death. But that does not justify convicting him of murder. Anyone who attempts murder, and goes about things as carefully and well as he can, is just lucky (or unlucky) if he does not cause the death he wishes to cause, but that does not warrant holding him for murder. So also for Simonson. In fact, he too attempted murder, and it is also just luck for him if he did not cause Summers' death.

It is arguable that if a man attempts murder, and it is just luck for him that he does not cause the death he wishes to cause, then morally speaking he has acted as badly as he would have acted had he succeeded.⁷ Many of those who take this view regard it as morally suspect that the penalty for murder should be heavier than the penalty for attempted murder. (Perhaps they view it as flatly unacceptable that there is such a difference in penalty. Perhaps they think the difference is just barely acceptable in light of the fact that imposing the same penalty might give unsuccessful attempters a motive to try again, or in light of some other, or additional, considerations.) At all events, the penalty for murder is everywhere heavier than the penalty for attempted murder. Or at least so I suppose; and we can anyway assume this true in the jurisdiction in which Tice is to be tried. So it is not enough to justify the charge of murder against Tice that it is just luck for him if he did not cause the death: to warrant imposition of the heavier penalty, the prosecution has positively to prove beyond a reasonable doubt *that* he caused it.

Well, isn't it beyond a reasonable doubt that Tice caused Summers' death? After all, it is .95 probable that he did.

III

It is often said that the kind of evidence available in *Smith v. Red Cab* and *People v. Tice* merely tells us the "mathematical chances"⁸ or the "quantitative

7. Thomas Nagel denies this claim in a very interesting article, *Moral Luck*, in T. NAGEL, *MORTAL QUESTIONS* 24-38 (1979).

8. *Smith v. Rapid Transit, Inc.*, 317 Mass. 469, 470, 58 N.E.2d 754, 755 (1945).

probability”⁹ of the defendant’s guilt. And it would be said that what is missing in those cases, the lack of which makes conviction suspect, is “real”¹⁰ or “individualized”¹¹ evidence against the defendant.

I strongly suspect that what people feel the lack of, and call individualized evidence, is evidence which is in an appropriate way causally connected with the (putative) fact that the defendant caused the harm.¹²

Consider the evidence that it was a Red Cab which caused the accident in *Smith v. Red Cab*: it consists entirely of Mrs. Smith’s testimony that a cab caused her accident, and that six out of ten of the cabs in town that night were operated by Red Cab. If we believe her, we believe there are such facts as that a cab caused her accident, and that six out of ten of the cabs in town that night were operated by Red Cab. But those facts lack an appropriate causal connection with the (putative) fact that Red Cab caused the accident.

What sort of causal connection would be appropriate? Well, if a witness came forward to say he saw the accident, and that the cab which caused the accident looked red to him, *then* we would have what would be called individualized evidence against Red Cab; and my suggestion is that that is because the accident-causing cab’s actually being red (and therefore being Red Cab’s) would causally explain its looking red to that witness. We might call this “backward-looking individualized evidence” of the defendant’s guilt because the bit of evidence (the witness’s believing the cab looked red to him) points back towards the (putative) fact that Red Cab caused the accident.

Or if it turned out that Red Cab had given a party for its drivers on the evening of the accident, a party which turned into a drunken brawl, then too we would have what would be called individualized evidence against Red Cab; and my suggestion is that that is because the party would causally explain its having been a Red Cab which caused the accident. We might call this “forward-looking individualized evidence” of the defendant’s guilt because the bit of evidence (the party) points forward towards the (putative) fact that Red Cab caused the accident.

Or more complicated (since it involves a common cause), if a red cab crashed into a parked car shortly after Mrs. Smith’s accident, and four blocks past the place of it, the driver giving all signs of being drunk, then that too would be called individualized evidence against Red Cab; and my suggestion is that that is because that driver’s having been drunk would causally explain *both* his crashing into the parked car *and* his (and therefore Red Cab’s) having caused Mrs. Smith’s accident.

9. *Day v. Boston & Me. R.R.*, 96 Me. 207, 217, 52 A. 771, 774 (1902).

10. *Id.* at 217, 52 A. at 774.

11. See, e.g., Tribe, *Trial by Mathematics: Precision and Ritual in the Legal Process*, 84 HARV. L. REV. 1329 (1971).

12. An earlier version of these ideas appears in section 8 of Thomson, *Remarks on Causation and Liability*, 13 PHIL. & PUB. AFF. 101 (1984). What I here call “individualized evidence,” in light of this term’s frequency in the literature, I there called “internal evidence.” Neither term strikes me as a particularly happy choice.

In the actual *Smith v. Red Cab*, no such further evidence came out. The facts available to us provide no forward-looking individualized evidence that Red Cab caused the accident, for they neither supply nor suggest any causal explanation of its having been a red cab which caused the accident. Moreover, the facts available to us neither supply nor suggest anything which might have been a common cause *both* of those facts available to us *and* of the (putative) fact that Red Cab caused the accident.

What is of interest is that we do have in the actual *Smith v. Red Cab* a piece of backward-looking individualized evidence for a hypothesis. Mrs. Smith says she could see it was a cab which came towards her, and its actually being a cab which came towards her would causally explain her believing this; so her saying she could see it was a cab which came towards her is backward-looking individualized evidence that it was a cab which caused her accident.

No one, of course, supposes that individualized evidence is (deductively valid) proof. In particular, our having backward-looking individualized evidence that a cab caused Mrs. Smith's accident is logically compatible with its not having been a cab which caused the accident.

Moreover, different bits of individualized evidence may differ in strength. For example, it is possible for a private car or bus or truck, or for all I know a gorilla, to be disguised as a cab, and the more non-cabs there are on the roads that are disguised as cabs, the less weight we are entitled to place on the causal hypothesis that Mrs. Smith's believing it was a cab which caused her accident was caused by its being a cab which caused her accident, and thus the less weight her believing it was a cab which caused her accident lends to the causal hypothesis that it was a cab which caused her accident. Still, her believing it was a cab which caused her accident is backward-looking individualized evidence that it was a cab which caused her accident, for her having that belief would be causally explained by its having been a cab which caused her accident.

Mrs. Smith's believing it was a cab which caused her accident would also be causally explained by its having been a red cab which caused her accident. (Again, her believing it was a cab which caused her accident would also be causally explained by its having been a cab once ridden in by a Presbyterian minister which caused her accident.) That does not mean that her believing it was a cab which caused her accident is backward-looking individualized evidence for the hypothesis that it was a red cab which caused her accident (or for the hypothesis that it was a cab once ridden in by a Presbyterian minister which caused her accident.) For there is no reason to think that the redness of a cab (or its past ridership) is causally relevant to its looking to a person like a cab. Mrs. Smith, of course, might be unusual in this respect: it might be that her retinas are so structured as to record "cabbiness" only when caused to do so by red-cabbiness. If we were given reason to think that that is true of her, we would thereby have been given reason to think her believing it was a cab which caused her accident *was* backward-looking individualized evidence that it was a red cab which caused her accident. But in the absence of a reason to

think her odd in some such respect as that, what we have is backward-looking individualized evidence only for the hypothesis that it was a cab which caused her accident.

For Red Cab to be guilty, the cab that came at Mrs. Smith (supposing it was a cab that came at her, as we do suppose if we believe her story) has to have had the features which distinguish Red Cab's cabs from the other cabs in town that night. Redness is one such feature, and no doubt there are indefinitely many others. But we have in hand no facts about the accident in which the (putative) redness of the accident-causing cab, or its (putative) possession of some other feature which distinguishes Red Cab's cabs, can be assigned an appropriate causal role. The facts available to us, then, provide (backward-looking) individualized evidence that a cab caused Mrs. Smith's accident, but no individualized evidence that the cab that caused the accident was one of Red Cab's cabs.

The point, then, is not that the only evidence we actually have in hand in *Smith v. Red Cab* is numerical or statistical,¹³ for we do have in that case a piece of individualized evidence for the hypothesis that it was a cab that caused the accident.

More important, numerical or statistical evidence too can be causally connected in an appropriate way with the (putative) fact it is presented to support. Suppose a plaintiff alleges that he was refused a job with a certain organization on grounds of race; in evidence, he presents statistics showing that the racial composition of the organization's work force diverges widely from that of the local population. Those data suggest a causal hypothesis, namely that the organization intends to discriminate in its hiring practices, and the organization's intending to discriminate in its hiring practices would causally explain *both* the existing divergence in racial composition, *and* the (putative) fact that it refused to hire the plaintiff on grounds of race. So that evidence too is individualized, although it is numerical or statistical.

If we had individualized evidence (and thus, on my hypothesis, appropriately causally connected evidence) against Red Cab, in addition to the evidence we already have in hand, then we would feel considerably less reluctant to impose liability on Red Cab. Why is that? That seems to me to be a very hard question to answer.

It cannot plausibly be said that the addition of individualized evidence against Red Cab would make us feel less reluctant to impose liability on it because the addition of individualized evidence against Red Cab would raise the probability that Red Cab caused the accident. Even in the absence of individualized evidence, the probability that Red Cab caused the accident is

13. It is of interest that although the court in *Smith v. Rapid Transit, Inc.* stated that "[t]he most that can be said of the evidence in the instant case is that perhaps the mathematical chances somewhat favor the proposition that a bus of the defendant caused the accident," 317 Mass. 469, 470, 58 N.E.2d 754, 755 (1945), the evidence in that case was entirely nonstatistical. I suggest that what the court felt the lack of was not nonstatistical evidence, but appropriately causally connected evidence.

already .6, which on a plausible interpretation of the requirements of tort law is higher than it need be.

Friends of the idea that individualized evidence is required for conviction have not really made it clear why this should be thought true. That has encouraged their enemies to suppose they have the idea because they think that individualized evidence is uniquely highly probabilifying.¹⁴ The enemies have then found it easy to make mincemeat of the friends. The enemies draw attention to the fact (and it is a fact) that eyewitness testimony, for example, which is a paradigm of individualized evidence, may be quite unreliable, that is, may probabilify to a lower degree than would some pieces of purely numerical or statistical evidence. And they draw attention to the mistakes about probability (and they are mistakes) which have been studied by Tversky and Kahneman, in particular, those which issue from ignoring base rates.¹⁵

But I think that is at best an ungenerous diagnosis of what is at work in the friends of individualized evidence. What is at work in them is not the thought that individualized evidence is uniquely highly probabilifying, but rather the feeling that it supplies something which nonindividualized evidence does not supply, which further something is not of value because it raises the probability of the hypothesis in question.

IV

I think it is helpful at this point to look at an analogous problem faced by the student of the theory of knowledge. What is currently called the classical, or traditional, account of knowledge says that a person A knows that a proposition p is true if and only if three conditions are met. First, p actually is true. Second, A believes that p is true. And third, A has a reason, and indeed, not just any old reason, but a good enough reason, for believing that p is true.¹⁶ (Some people prefer to express this third condition in the words "A is

14. Some friends of individualized evidence do seem to think it uniquely highly probabilifying. Consider, for example, the court's comment in *Day v. Boston & Me. R.R.*:

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 sixes will fall uppermost 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.

96 Me. 207, 217, 52 A. 771, 774 (1902). But the passage is obscure, and perhaps not to be understood in the way that at first suggests itself. In any case, I take it that friends of individualized evidence who would express their view in words such as these are in a very small minority.

15. For a very helpful survey and further references, see Saks & Kidd, *Human Information Processing and Adjudication: Trial by Heuristics*, 15 LAW AND SOC'Y REV. 123 (1980-81). Judging from that article, I think that Saks and Kidd would say, not only that the friends of individualized evidence are mistaken in thinking it of any special interest, but also that it is a mistake to suppose there is a problem to be solved here—thus that it is a mistake to think it would be worrisome for Mrs. Smith to win her case in *Smith v. Red Cab* or for the people to win its case in *People v. Tice*.

16. It is arguable that we should add: "and it is because of having a good enough reason for believing that p is true that A does believe that p is true." (For it is arguable that a man might have a good enough reason for believing that p is true, and yet believe that p is true for no reason at all, or for a bad reason.)

justified in believing that *p* is true.” Others express it in the words “A has adequate evidence for *p*.”) Not surprisingly, it is this third condition which has generated the controversy. Nobody argues about the first or second conditions: it is plain enough that you know that the sun is shining only if it is shining and you believe it is. What is unclear is exactly what is required for A to have a reason which is good enough for it to be true that A satisfies the third condition.

But one thing is clear: if A’s satisfaction of those three conditions is supposed to be sufficient for A’s knowing that *p* is true, then A is not marked as satisfying the third condition by virtue of its being the case that in light of all of the facts available to A, he is entitled to conclude, indeed is rationally required to conclude, that *p* is highly probable. Suppose Alfred believes, truly, that Bert bought five tickets in a lottery, and that a hundred tickets were sold altogether. Suppose also that those are the only facts such that Alfred is both aware of them and aware of their bearing on the question whether Bert will win the lottery. (For example, Alfred is not aware of any reason to believe that the lottery is rigged, or, if he is aware of a reason to believe it rigged, there is no ticket such that he is aware of a reason to believe the lottery is rigged in favor of that ticket.) The following is therefore true:

In light of all of the facts available to Alfred, he is entitled to conclude, indeed is rationally required to conclude, that it is .95 probable that Bert will lose the lottery.

Alfred is aware that he is so situated, and for that reason concludes that Bert will lose the lottery, so that it is also true that:

Alfred believes that Bert will lose the lottery.

Suppose, lastly, that Bert will actually lose the lottery (though of course Alfred does not, and we may suppose we ourselves do not, find this out until later). Then it is also true that:

Bert will lose the lottery.¹⁷

But it is plainly *not* true that:

Alfred knows that Bert will lose the lottery.

It pays to stress: not merely is it not true to say now that Alfred knows that Bert will lose the lottery, it will also not be true to say later, after Bert has already lost the lottery, that Alfred then knew (that is, in advance) that Bert would lose the lottery.

There is something missing in Alfred, something the lack of which makes it false to say he knows that Bert will lose. The point obviously is not that .95 is not a sufficiently high probability. Pick any probability you like, as high as you like, and we need only suppose that the appropriate number of tickets were

17. Some people think that sentences which are contingent and in the future tense have as yet no truth-value. If that were true, then it would be wrong to say we speak truly if we say “Bert will lose the lottery,” even if Bert does later lose the lottery. I say “if that were true” because I think it false. No matter. Anyone who holds this view is invited to substitute for the lottery a game in which Bert just drew five cards at random from a Superdeck (which contains 100 successively numbered cards), and then replaced them without looking. Alfred believes none of Bert’s cards said “73.” Let us suppose that is true; and so on, and so on.

sold to entitle Alfred to conclude it probable in that degree that Bert will lose the lottery. Still, even if in fact Bert will lose, Alfred does not know that he will.

Should we say that what is missing in Alfred is something which would make him satisfy a further, fourth, condition on knowledge? That is, it could be said that Alfred's reason for believing Bert will lose the lottery is good enough for him to satisfy the third condition, and that what the case shows is that knowledge requires something more than satisfaction of the three conditions laid out in the classical account. Or should we instead say that although Alfred has good reason for believing Bert will lose the lottery, what is missing is something which would make his reason be *good enough* for him to satisfy the third condition on knowledge? It makes no theoretical difference which choice we make, for the question what is missing in Alfred remains to be answered whichever choice is made. But I prefer the second, for the classical account of knowledge is so natural and intuitively plausible as to incline one to want to interpret it in such a way as to have it be immune to such cases as that of Alfred and Bert.

To get at exactly what is missing in Alfred, then, one would have (as I shall put it) to get at exactly what is required for A to have a reason which is good enough for it to be true that A satisfies the third condition on knowledge. That, alas, is a hard problem.

But it is easy enough to say very roughly what is missing in Alfred: what he lacks is something which would make it not be just luck for him that Bert will lose the lottery.

This is a helpful way of expressing what Alfred lacks¹⁸ because it points towards anyway one of the things which is required for A to have a reason that is good enough for it to be true that A satisfies the third condition on knowledge. What I have in mind is this: it seems very plausible to think that A's reason for believing that p is true must ensure, or *guarantee*, that p is true. Alfred's reason for believing that Bert will lose the lottery does not in any way guarantee that Bert will lose it. By contrast, consider Alice. Alice believes that Bertha bought one ticket in a certain lottery, and that unbeknownst to Bertha, the ticket seller tore up her ticket stub directly after selling it to her, and that the drawing will be made from among the stubs. Suppose that belief of Alice's is true. Then Alice has good reason for believing that Bertha will lose the lottery, namely the fact that Bertha bought one ticket, the ticket seller tore up her stub, and the drawing will be made from among the stubs. Suppose Alice does therefore believe that Bertha will lose the lottery. Suppose, lastly, that time passes, and that it has now turned out that Bertha did lose the lottery, and that her losing it *was caused by* the ticket seller's tearing up her stub. Then the fact which was Alice's reason for believing that Bertha would lose the lottery caused, and thereby guaranteed, that Bertha would lose it. This makes it seem much less implausible to think Alice knew

18. It was suggested by Thomas Nagel's remarks on epistemology in T. NAGEL, *supra* note 7.

that Bertha would lose the lottery than it did to think Alfred knew that Bert would lose it.

Again, consider Arthur. Arthur believes he is having, and is in fact having, a visual impression as of a chicken in front of him, and his having that visual impression is his reason for believing there is a chicken in front of him. If it later turns out that there was then a chicken in front of him, and that there being a chicken in front of him *was causally necessary* for his having that visual impression, then his having that visual impression guaranteed that there was a chicken in front of him. This makes it seem much less implausible to think Arthur knew that there was a chicken in front of him than it does to think Alfred knew that Bert would lose his lottery. It is similarly so for more complicated cases, in which the guaranteeing proceeds via common causes.¹⁹

On most views about knowledge, even this is not enough. Suppose that A believes there is such a fact as *q*, and takes *q* to be reason for believing that *p* is true. On most views about knowledge, it is not enough to secure A's satisfying the third condition on knowledge that there actually is such a fact as *q*, and that *q* actually guarantees that *p* is true: in addition (on those views) A must have good reason for believing that there is such a fact as *q*, and A must believe, and have good reason for believing, that *q* does guarantee that *p* is true. Or perhaps, more strongly, that *these* good reasons must themselves contain guarantees. On some views about knowledge, if there is luck anywhere at all in A's route to his belief that *p* is true, then it is just luck for A if *p* is true, and A therefore does not know that *p* is true.

Those (most restrictive) views about knowledge had better be false, for down this road lies skepticism, that is, the thesis that nobody knows anything. (Only people who are excessively charmed by philosophy can regard it as a happy outcome that nobody so much as knows he has hands or feet.)

I think it helpful at this point to turn from knowledge itself to saying one has it. That in any case is what will matter for us.

If A says to B "I know that *p* is true," then A does something more than just assert something: he gives B his word that *p* is true, and in one of the strongest ways we have of doing this. Indeed, A positively invites B to take his word for the truth of *p*, and to rely on it. That being so, and because it is so, there are rather strong moral constraints governing the acceptability of saying "I know that *p* is true." Plainly A's saying this is not made acceptable by the mere fact that *p* turns out to be true. Suppose Boris will suffer a loss if he relies on the truth of *p*, where *p* is not true. He trusts Andrew; he tells Andrew he will suffer a loss if he relies on the truth of *p* where *p* is not true, and asks Andrew if Andrew knows whether *p* is true. Andrew says "Yes, I know that *p* is true," but he says this in awareness that he has not the slightest idea whether *p* is true—he does so because he decided to flip a coin, heads I

19. The guaranteeing required for knowledge need not proceed via causes at all. Compare mathematical knowledge: the truth of one statement may guarantee the truth of another via entailing it.

say it, tails I don't, and the coin came up heads. Then Andrew has acted improperly in saying what he said, even if, as it turns out, *p* is true.

One way, the weaker way, of explaining the impropriety in Andrew is to point to the fact that he said "I know that *p* is true" without having any reason at all, and a fortiori without having good reason, to think *p* true. Boris had said he would suffer a loss if he relied on the truth of *p* where *p* is not true; lacking good reason to think *p* true, Andrew could not be sure that *p* was true, and thus could not be sure that Boris would not suffer that loss if he relied on the truth of *p*—yet he all the same, and therefore wrongly, invited Boris to rely on it.

But that is weak. Consider Alfred again. Alfred has good reason to think Bert will lose the lottery he entered, for Alfred is entitled to believe it .95 probable that Bert will lose the lottery. Suppose the lottery works like this, and Alfred knows it does: the winner must prove he is the winner by producing his ticket, or he must wait six months for his prize, during which time the lottery organizers will assure themselves that he is who he says he is. Nevertheless Alfred says to Bert: "Look, it's silly of you to hang on to those five, by now grubby, lottery tickets. I *know* you are going to lose that lottery." Bert will suffer a loss (six months' interest on a large sum of money) if he accepts Alfred's invitation to rely on his losing the lottery, throwing out his tickets, and it then turns out that he wins the lottery; and Alfred knows this. The fact that Alfred is sure that Bert will lose the lottery, and thus that Bert will not suffer that loss, does not make it acceptable for Alfred to have said what he said.

Let us call the loss a person will suffer if he relies on the truth of *p* where it turns out that *p* is not true that person's "potential mistake-loss." Then we can say the fact that Alfred is sure that Bert will lose the lottery, and thus that Bert will not suffer his potential mistake-loss, does not justify Alfred's saying to Bert, "I know you will lose the lottery."

What Alfred implies by saying "I know you will lose" is not merely that he is sure that Bert will lose, and thus that Bert will not suffer his potential mistake-loss, but that Alfred has a certain kind of ground for being sure of this—"insider's information" (about rigging, perhaps) in the case of lotteries, stock markets, and other gambler's games. Alfred implies he has a guarantee of some kind that Bert will lose, so that not merely will Bert not suffer his potential mistake-loss if he relies on Alfred's word, but that Bert does not even risk suffering his potential mistake-loss if he relies on Alfred's word. (Think of Bert's reaction if he accepts Alfred's invitation to rely on his losing, and tears up his tickets, and then later learns that the ground on which Alfred issued his invitation was merely the large number of tickets sold, which made it highly probable that Bert would lose. Bert himself knew it was highly probable he would lose at the time of buying his tickets; and he will correctly view himself as having no less right to complain about Alfred's behavior if he loses the lottery than if he wins it.)

It is plausible to think, quite generally, then, that if A is aware that B will suffer a loss if he relies on the truth of *p* where *p* is not true, then A ought not say to B "I know that *p* is true" unless A is more or less sure that he has a guarantee that *p* is true²⁰—so that B not only will not suffer, but does not even risk suffering, his potential mistake-loss. There are, after all, considerably weaker sentences than "I know that *p* is true" which the language makes available to A: he can say instead "I believe that *p* is true" or "It is highly probable that *p* is true" or "I believe that *p* is true because it is highly probable that *p* is true," and so on. To assert some such weaker sentence as this is not to do nothing: assertions of them are governed by moral constraints, too. But none of them is such that to assert it is to invite one's intended hearer to take one's word for the truth of *p* and rely on it. Doing *that* calls for being more or less sure that one has a guarantee of the truth of *p*—at any rate, it does where a potential mistake-loss is in the offing.

By way of making connection with what was said earlier, we could restate this point as follows: doing *that* calls for being more or less sure that it would not be just luck for one if *p* turns out to be true—at any rate, it does where a potential mistake-loss is in the offing.

This gives us a second, and stronger, way of explaining the impropriety in Andrew. He told Boris he knew that *p* was true because the coin he flipped came up heads. This was wrongful behavior on his part because he invited reliance without being in the slightest degree sure of having a guarantee, despite his awareness of Boris's potential mistake-loss.

A need only be, as I said, "more or less sure" that he has a guarantee, for circumstances differ. Suppose you ask me now, in the afternoon, if there's cold chicken in the icebox. I say, "Yes, I'm certain there is, in fact I know there is." Why do I say that? Because I believe I put some there this morning, and believe my putting some there this morning guarantees that there is some there now. In the normal course of events, my saying these words is acceptable, even if it turns out that there is no cold chicken in the icebox now, for example, because some burglars broke in at noon and ate it while burgling. That is because in the normal course of events I know your potential mistake-loss is small—you will merely waste a bit of energy on a trip to the icebox, and suffer a minor disappointment, if it is "chickenless." But we could imagine a series of cases in which your potential mistake-loss increases; that would be a series of cases in which the burden on me to make sure that I really do have a guarantee increases proportionately. Where I know your potential mistake-loss is truly terrible, then the burden on me is

20. Usage is slippery, and there are cases which suggest that this quite general claim may be overbroad. Suppose Alfred says to Bertha, "I know that Bert will not pay you the money he owes you," and says this on the ground that Bert has no money other than that which he will get if he wins the lottery, and that a large number of tickets were sold; wouldn't *that* be all right? (I owe this example to David Gauthier.) I strongly suspect, however, that this seems to us acceptable (if it does) only because we are not imagining the further details which would be required to make plain what Bertha's potential mistake-loss is; and I predict that the supplying of such further details would make Alfred's behavior seem unacceptable.

very heavy indeed. Did my putting cold chicken in the icebox this morning really guarantee that there is some there now? Your potential mistake-loss being so great, I ought not ignore the possibility that it did not. (Burglars are only one of many more or less weird possibilities.) I had better not invite you to run the risk of suffering that truly terrible loss unless I make *very* sure I have a guarantee—as, for example, by getting up out of my chair and going to have a look, and perhaps also a sniff and taste.

In light of these considerations it might well be wondered why we ever say anything so strong as “I know that *p* is true.” Well, we do not normally say this unless we believe our hearer will gain something if he acts on the supposition that *p* is true where it is true, or suffer a loss if he fails to act on the supposition that *p* is true where it is true. (In the normal course of events, a person who asks whether there is cold chicken in the icebox wants to eat some, and thus has something to gain by acting on the supposition that there is some there if it turns out that there is). A foregone gain is what is sometimes called an “opportunity cost,” and thus itself a kind of loss, so it could have been said, instead, more briefly: we do not normally say, “I know that *p* is true,” unless we believe our hearer will suffer a loss (if only a loss which is an opportunity cost) if he fails to act on the supposition that *p* is true where it is true. What would be the point of inviting a hearer to rely on the truth of *p* if the hearer has nothing at all to lose by failing to rely on it even if *p* is true?

There might, then, be a case in which A believes both of the following about B. First, B will suffer a loss if he acts on the supposition that *p* is true where it is not true—that is B’s potential mistake-loss. Second, B will suffer a loss if he fails to act on the supposition that *p* is true where it is true; I will call this B’s “potential omission-loss.” Here A has some more or less delicate balancing to do. But perhaps this much is clear enough: the greater the amount by which B’s potential mistake-loss exceeds B’s potential omission-loss, the more sure A must be that he has a guarantee that *p* is true before saying, “I know that *p* is true.”

It is not knowledge itself, but rather saying one has knowledge, that will matter for us. But I think it is worth indicating, just briefly, why (as I said) I think it helps a student of knowledge to attend to what is required for acceptably saying one has it. According to the classical account of knowledge, A knows that *p* is true if and only if *p* is true, A believes that *p* is true, and A has good enough reason for believing that *p* is true. What marks A as having good enough reason for believing that *p* is true? I suggest it is enough if A takes himself to have a guarantee that *p* is true, and if, also, what he takes to guarantee this does guarantee it. Thus, for example, suppose (again) that I believe there is cold chicken in the icebox because I believe I put some there this morning, and believe my putting some there this morning guarantees that there is some there now. Suppose also that my believing I put some there this morning guarantees that I did put some there this morning, and that that guarantees that there is some there now. Then I do know that there is some

there now. I could have been wrong: for example, my putting cold chicken in the icebox this morning could have failed to guarantee that there is some there now. (Perhaps burglars took the chicken.) I suggest we should take that fact to have a bearing, not on whether I know, but only on what I may acceptably say.

V

I fancy it may already be clear from this long digression on knowledge what, on my view of them, is at work in the friends of individualized evidence.

The jurors do not say, “we know that the defendant is guilty” at the close of the trial; they say only, “the defendant is guilty”; but in saying that they do something of great significance. It is not strong enough to say they declare the defendant guilty. If you and I have been watching the trial, I may say to you as we leave the courtroom, “the defendant is guilty.” I have declared the defendant guilty, but have not done, because I am not so situated as to be able to do, what the jurors do when they say these words at the close of the trial. The institution in which they are participating is so structured that their saying these words then is their imposing liability on the defendant—for if they say these words at that time, appropriate others will act on the supposition that he is guilty, which includes imposing the relevant penalty on him. So they do not merely invite reliance, they act in awareness that reliance will follow.

Under what conditions is it acceptable for the jurors to agree to say those words at the close of the trial? One thing which is perfectly plain is that their agreeing to say those words is not made acceptable by the mere fact that the defendant actually is guilty of what he is charged with. That what the jurors declare true turns out to have in fact been true does not by itself make it acceptable for them to have declared it true.

This point is obvious, but it pays to make its source explicit. Suppose that a jury is puzzled by the evidence which has been presented to it, and cannot arrive at a consensus as to its weight. “I know,” says one juror, “let’s decide by flipping a coin—heads we impose liability, tails we don’t.” They agree; they flip a coin, which comes up heads; so they return and say, “the defendant is guilty.” Their doing that is not made acceptable by the fact (supposing it a fact) that the defendant actually is guilty. If the defendant is guilty, then he deserves the penalty which this jury causes to be imposed on him; but that the defendant not suffer the relevant penalty unjustly is not all that matters to us. It matters to us, not just that a defendant not suffer a penalty unjustly, but also that the penalty not be imposed on him unjustly.

The defendant will suffer the penalty unjustly if he is not guilty, and so does not deserve the penalty; that means that it is unjust to impose liability on him, and thereby cause him to suffer the penalty, unless one believes one has good reason to believe that he is guilty, and therefore deserves the penalty. That being so, we can say, and we have an explanation of why we can say, that the jury I just described imposed liability unjustly: they imposed liability

without believing they had good reason to believe that the defendant was guilty.

There is a second, and stronger, possible explanation of why we can say that that jury imposed liability unjustly: it was just luck for those jurors if what they declared true was true—just luck for them if it actually was the case that the defendant was guilty.

That is stronger. Consider the jury in the hypothetical case I called *People v. Tice* in section II above. Suppose it declares Tice guilty of murder, not on the ground that a coin was tossed and came up heads, but on the following two grounds. First, the evidence makes clear (perhaps Tice has even confessed) that he attempted to kill Summers. Second, Summers was killed by one of one hundred pellets fired at him, and ninety-five of the pellets fired at him were fired by Tice. Then the jury imposes liability on Tice on the ground of what is on any view good reason to believe Tice guilty. (Its situation in respect of Tice's being guilty is exactly like Alfred's situation in respect of Bert's losing the lottery: in both cases, rationality requires believing the conclusion highly probable.) If it is required of a jury only that it not impose liability without good reason to believe the defendant guilty, then this jury does not impose liability unjustly. All the same, it is just luck for the jury if it actually was Tice who killed Summers, and thus if Tice committed murder. So if it is required of a jury that it not impose liability unless it has, not merely good reason, but reason of a kind which would make it not be just luck for the jury if its verdict is true, then this jury imposes liability unjustly.

On my view of them, what is at work in the friends of individualized evidence is precisely the feeling that just imposition of liability requires that this stronger requirement be met.²¹ They believe, as they say, 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. I say "not without reason" because I feel in considerable sympathy with them.

What would make it not be just luck for the jury if what it declares true is true? A guarantee. I suggested that individualized evidence for a defendant's guilt is evidence which is in an appropriate way causally connected with the (putative) fact that the defendant is guilty, and hence (putatively) guarantees the defendant's guilt; so to require individualized evidence of guilt just is to be requiring a guarantee.

None of this is incompatible with the fact that there is a difference made in our law between the standard of proof required in criminal cases on the one

21. To say this is to rest the importance of individualized evidence on what the jurors do, rather than on something external to the trial. Compare Laurence H. Tribe, who rests it on the symbolic or expressive function of trials. See Tribe, *supra* note 11. Compare also Charles Nesson, who rests it on the messages communicated by trials. See Nesson, *The Evidence or the Event? On Judicial Proof and the Acceptability of Verdicts*, 98 HARV. L. REV. 1357 (1985).

hand, and cases in tort on the other hand. Our law requires the jury in a criminal case to be sure beyond a reasonable doubt that the defendant is guilty before imposing liability on him; the friend of individualized evidence may be taken to say that the jury must be sure beyond a reasonable doubt that the defendant is guilty *because of* being sure beyond a reasonable doubt that there are facts available to it which guarantee that the defendant is guilty. Our law requires the jury in a case in tort to believe no more than that it is more probable than not that the defendant is guilty; the friend of individualized evidence may be taken to say that the jury must believe it is more probable than not that the defendant is guilty *because of* believing it more probable than not that there are facts available to it which guarantee that the defendant is guilty.

We met an analogous difference in section IV above, and the differences have analogous sources. Our society takes the view that, in a criminal case, the loss to society if the defendant suffers the penalty for a crime he did not commit is very much greater than the loss to society if the defendant does not suffer the penalty for a crime he did commit. (That is, in particular, we think it would be considerably worse that an innocent defendant be punished than that a guilty defendant go free.) This point might be reexpressed as follows: our society takes the view that in a criminal case, the society's potential mistake-loss is very much greater than the society's potential omission-loss. It would be no wonder, then, if our law imposed a heavy standard of proof on the jury in a criminal case; and according to the friend of individualized evidence, that means the jury must be very sure of having a guarantee before imposing liability for a crime. The fact that the standard of proof in a case in tort is more relaxed by itself suggests that our society takes the view that, in a case in tort, the society's potential mistake-loss is not much greater than the society's potential omission-loss. (Thus, in particular, we think it would be worse that an innocent defendant be forced to pay a plaintiff who either was not really harmed at all, or who anyway was not harmed by the defendant, than that a guilty defendant go free and the plaintiff he injured go uncompensated—but not by much.)

VI

The typical hypothetical cases which descend from *People v. Collins* raise the problem we are dealing with less cleanly than *People v. Tice*, and that in two ways.

Here is an example. Mrs. Smith testifies that she saw a man, indeed exactly one man, kill Bloggs, and that she could see he was one-legged, left-handed, entirely bald, and extremely tall. Mrs. Jones is a biologist-statistician, and she testifies that men with all four of those features are very rare: only one man in ten million has all four of them. There is the defendant Mullins, and we can see that *he* has all four of them. So we may be inclined to think we may conclude, on the basis of that evidence alone, and without hearing or seeing anything more, that it is highly probable that Mullins killed Bloggs.

As the commentators on this kind of case enjoy pointing out, however, this would be a mistake on our part. After all, there are a lot more than ten million men in the world. How many are there? Let us suppose there are a billion men in the world, thus 1,000 million. If, as Mrs. Jones' data suggest, one in ten million men have all four features, then 100 in 1,000 million have all four features; so we are entitled to conclude only that the probability that Mullins killed Bloggs is one in 100, thus .01—not at all a high probability.

I think it helps to see this point if we lay our information out in the way in which I laid our information out in *People v. Tice* and *Smith v. Red Cab*. Thus we believe, on the basis of what we can see on looking at the defendant Mullins:

(1'') Mullins is a man who is one-legged, left-handed, entirely bald, and extremely tall.

If we believe Mrs. Smith's story, we believe that exactly one man killed Bloggs, and that the man is one-legged, left-handed, completely bald, and extremely tall. If we also believe Mrs. Jones's story, we believe that one in ten million men have those four features. If we believe there are 1,000 million men, we are committed to believing that 100 men all told have those four features. So we are committed to believing:

(2'') One of the one hundred men who are one-legged, left-handed, entirely bald, and extremely tall killed Bloggs.

The conclusion we are entitled to draw is only:

(3'') The probability that Mullins killed Bloggs is .01.

Our problem is this: what should we think of the idea of convicting a defendant on the basis of evidence which makes it highly probable he is guilty, but which is nonindividualized evidence against him? So the first way in which the typical hypothetical descendant of *People v. Collins* is less clean than *People v. Tice* is this: in that kind of case, the available data do not even entitle us to conclude that it is highly probable that the defendant is guilty.

This may seem to be a relatively trivial difficulty, for there are a number of different ways in which the case against Mullins could be revised so as to raise the probability of his guilt, despite its anyway seeming to remain a case in which no individualized evidence is produced against him. We might suppose, for example, that Mrs. Smith testifies that the man she saw kill Bloggs had, not merely the four features I have mentioned, but in addition, only one eye—not one good eye and one bad eye, but exactly one eye, right in the middle of his forehead. Mullins, as we can see in the revised case, has not merely the four features, but exactly one eye, right in the middle of his forehead. Our biologist-statistician, Mrs. Jones, assures us that she would have said this was impossible if she had not seen Mullins: it is due to a freak genetic mutation which may be expected to occur no more often than once in 100 billion men. There being in existence only one billion men, it is highly probable that Mullins is the only existing man with all five features, and thus (if we believe Mrs. Smith's story) it is highly probable that Mullins did kill Bloggs.

This kind of move has to be watched, however, for once the probability is made high in this way, our problem does not come out more cleanly; it instead disappears. What I have in mind is this. In the original case against Mullins, Mrs. Smith believes she saw a one-legged, left-handed, entirely bald, and extremely tall man kill Bloggs. That is individualized evidence that a man with those four features killed Bloggs, for the (putative) fact that a man with those four features killed Bloggs would causally explain Mrs. Smith's believing she saw a man with those four features kill Bloggs. (Compare the fact that in *Smith v. Red Cab* Mrs. Smith's believing she saw a cab come towards her was individualized evidence that it was a cab that came towards her.) Our further evidence also suggests there are other men than Mullins, ninety-nine other men in fact, who have all four features, so getting individualized evidence against Mullins requires getting some fact in respect of which an appropriate causal role is played by a feature which distinguishes Mullins from the others. (Compare the need of a fact in respect of which an appropriate causal role is played by a feature which distinguishes Red Cab's cabs from the other cabs we believe were in town on the night of the accident in *Smith v. Red Cab*.)

In the revised version of the case, Mrs. Smith believes she saw a one-legged, left-handed, entirely bald, extremely tall, *and* one-eyed man kill Bloggs. That is individualized evidence that a man with those five features killed Bloggs, for the (putative) fact that a man with those five features killed Bloggs would causally explain Mrs. Smith's believing that she saw a man with those five features kill Bloggs. But *here* our further evidence suggests that only Mullins has all five features, and therefore that there is no such thing as a feature which distinguishes him from the other men who have all five features, and therefore that there is no possible fact in respect of which such a distinguishing feature so much as could play an appropriate causal role. To the extent to which we believe that further evidence, then, we shall take ourselves to have individualized evidence, not merely that a man with those five features killed Bloggs, but that Mullins did—he being the only available candidate with the five features.

There is more to be said here, but I bypass it, because whatever revisions might be made in the case against Mullins, there remains the second way in which the typical hypothetical descendant of *People v. Collins* raises our problem less cleanly than *People v. Tice*: in that kind of case, the sources of the statistical information are typically softer than the source of the piece of statistical information in *People v. Tice*. For example, that exactly one man in ten million has all four of the features which interested us in the original case against Mullins may be true; but we may well be pardoned if we do not feel much confidence in those figures. By contrast, I was inviting you to suppose that the evidence presented in court entitles us to feel entirely confident that exactly one hundred pellets were fired at Summers, and that exactly ninety-five of them were fired at him by Tice. (And compare *Smith v. Red Cab*. Perhaps we do not feel *entirely* confident that six out of ten of the cabs in town

that night were operated by Red Cab, since we are so familiar with gypsy cabs. But we may have good reason to believe in the figures Mrs. Smith gives us, and no positive reason to disbelieve in them, perhaps in that Red Cab itself gives us no reason to think them wrong.)

On the other hand, I am not arguing that no hypothetical descendant of *People v. Collins* raises our problem cleanly. Quite to the contrary: *People v. Tice* itself descends from it.

VII

Commentators on this kind of case very commonly say that all evidence is "ultimately" statistical or probabilistic,²² and that this itself means there is muddle in the friends of individualized evidence. Such comments are decidedly not transparent, and they can be interpreted in a number of different ways. The interpretation which makes them most plausible, however, seems to me to be one under which they do not really constitute a difficulty for the friends of individualized evidence.

Under that interpretation, to say all evidence is ultimately statistical says about causal hypotheses in particular that such evidence as we have for them is itself statistical (or anyway, is ultimately itself statistical), so that we can have no more confidence in the truth of a causal hypothesis than we have in the statistical data which (ultimately) support it. This does not strike me as obvious, but perhaps it is true. Perhaps it is true that I am entitled to place no more confidence in the hypothesis that my pressing the button marked "T" just now caused a "t" to appear on the page than I have in some set of statistical data. Alas, I do not know what exactly the members of the set are, but perhaps one member is not quite that every time I press the button marked "T" a "t" appears on the page, for sometimes I have previously pressed the button marked "SHIFT," so that a "T" appeared on the page instead, and sometimes the machine had not been turned on, so that nothing appeared on the page, and so on, but something *like* this.

Now why might it be thought that this should trouble the friend of individualized evidence? He is familiar with the idea that a causal hypothesis may be supported by statistical data. Compare, for example, the kind of race discrimination case mentioned in section III above, in which statistical data about the divergence between the racial composition of an organization's workforce and that of the local population suggests the causal hypothesis that the divergence was caused by an intention to discriminate in hiring.

Moreover, he is familiar with the fact that being presented with statistical data may entitle us to feel more or less confidence in the truth of a causal hypothesis. As I said in connection with *Smith v. Red Cab*, the more non-cabs there are on the roads that are disguised as cabs, the less confidence we are

22. Tribe himself says, "I am, of course, aware that *all* factual evidence is ultimately 'statistical,' and all legal proof ultimately 'probabilistic'." Tribe, *supra* note 11, at 1330 n.2. See also Saks & Kidd, *supra* note 15, at 153 ("Invariably, all information is really probability information.").

entitled to place in the hypothesis that its being a cab that came towards Mrs. Smith does causally explain her believing she saw a cab come towards her. Similarly for a surprise witness, who comes forward to say he saw the accident, and that the accident-causing cab looked red to him: the more often non-red things look red to that witness when he is placed in similar circumstances, the less confidence we are entitled to place in the hypothesis that the cab's being red does causally explain his believing that it looked red to him.

One can, after all, be more or less sure of having the kind of guarantee that, as I said, is what the friends of individualized evidence think some degree of assurance is necessary for just imposition of liability. That our assurance of having a guarantee of the appropriate kind rests (ultimately) on statistical data seems to me to be something he can in consistency agree to.

And it need not trouble the friend of individualized evidence that the particular causal hypothesis that this or that bit of individualized evidence points to may not be well supported by the available statistical data, and hence that the bit of individualized evidence may not lend much weight to the hypothesis that the defendant is guilty. As I said, it is an ungenerous diagnosis of what is at work in the friends of individualized evidence to take them to think it of value because of thinking it uniquely highly probabilifying. What interests them is something else—a something else that I have been trying to bring out, and to invite others to take an interest in.

