

Epistemic Sensitivity and Evidence

Mario Günther*

Abstract

In this paper, we put forth an analysis of sensitivity which aims to discern individual from merely statistical evidence. We argue that sensitivity is not to be understood as a factive concept, but as a purely epistemic one. Our resulting analysis of epistemic sensitivity gives rise to an account of legal proof on which a defendant is only found liable based on epistemically sensitive evidence.

Keywords. Legal Epistemology, Sensitivity, Statistical vs. Individual Evidence

1 Introduction

In a civil lawsuit, a plaintiff presents evidence in order to convince a fact finder that a defendant is liable for some wrongful act. The fact finder – a judge or a jury – rules whether or not the defendant is liable for the plaintiff's claim. The verdict of the fact finder is governed by a burden of proof. Legal proof in civil law requires the prosecution to establish the defendant's liability on the *preponderance of the evidence*. This burden of proof is typically interpreted thus: a plaintiff's claim counts as proven in court just in case the claim is established to be more likely than not after considering the available and admissible evidence. The preponderance of the evidence standard gives rise to a simple account of legal proof: a defendant should be found liable if and only if (iff) the probability that the defendant is guilty surpasses $1/2$ given the available and admissible evidence.

*Mario.Guenther@anu.edu.au

Let us spell out this evidential standard. Where P is a probability function, p a proposition expressing a claim, and e a proposition expressing the available and admissible evidence, we may say that the claim p meets preponderance of the evidence iff $P(p \mid e) > 1/2$. To illustrate how this burden of proof is meant to apply, consider the following scenario:

Witness. The organizers of the local rodeo decide to sue John for gatecrashing their Sunday afternoon event. All the available and admissible evidence is as follows: John attended the Sunday afternoon event – he was seen and photographed on the main ranks. No tickets were issued at the entrance, so John cannot be expected to prove that he bought a ticket with a ticket stub. A witness observed John climbing the fence and taking a seat. There is strong and uncontested evidence that the fairly reliable witness testimony makes it 70% likely that John gate-crashed.¹

Many authors claim that a fact finder will find John liable for gatecrashing in *Witness*. However, according to the same authors, there are similar cases where the fact finders will not find John liable, or at least not obviously so.² Consider an example inspired by Cohen (1977, pp. 74-81):

Paid per Person. The organizers of the local rodeo decide to sue John for gatecrashing their Sunday afternoon event. All the available and admissible 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 at the entrance, 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.

¹This fictional example and the one to come are based on the examples due to Blome-Tillmann (2015).

²These authors include Enoch et al. (2012, p. 198), Blome-Tillmann (2015, p. 103), Gardiner (2018, p. 181), and Di Bello (2019, pp. 1045-6, 1080). Ross (2020), by contrast, argues that there are *some* cases, where fact finders should find liable based on statistical evidence.

Why would a fact finder not find John liable for gatecrashing in *Paid per Person*? The burden of proof seems to be met. The probability that he gatecrashed given the available and admissible evidence is $7/10$. Still it does not seem right to find John liable based on the mere statistics. He was just too randomly picked out. It seems as if imposing liability on John merely based on the statistical evidence that 70% of attendees gatecrashed is unwarranted and unjust. Rather we would need non-statistical or ‘individual’ evidence, as in *Witness*. The testimony of an eye-witness is intuitively more specific to John’s behavior, and it seems just to impose liability based on this individual evidence.

In the face of cases like *Paid per Person*, many have argued that legal burdens of proof should not be understood merely in terms of probabilities.³ The argument runs roughly as follows: the probability of John gatecrashing given the available and admissible evidence exceeds $1/2$ in both cases. Indeed, the probability of John gatecrashing is equal by design. Hence, in order to account for the intuitive difference between the cases, a burden of proof requires a condition in addition to meeting a certain probabilistic threshold. But what condition?

Enoch et al. (2012) discuss the condition that legal proof requires *sensitive* evidence. The idea is that the eye-witness testimony is sensitive because, if John had not gatecrashed, we would not have believed the testimony that John climbed the fence. By contrast, the 700:300 statistics is not sensitive. If John had not gatecrashed, we would still have believed that 700 out of 1000 attendees gatecrashed. They generalise this idea: individual evidence is sensitive, while merely statistical evidence is not. The additional condition that legal proof requires sensitive evidence faces problems, but we think it is worth exploring.

In Section 2, we explain Enoch et al.’s analysis of sensitive evidence and two of its problems. We then propose an analysis of epistemic sensitivity in Section 3. Our analysis overcomes the two problems of Enoch et al.’s and similar analyses, and gives rise to an account of legal proof. Finally, in Section 4, we argue that our analysis of epistemic sensitivity can be seen as

³For example, Cohen (1977), Enoch et al. (2012), Enoch and Fisher (2015), Haack (2014), Smith (2010, 2018, 2020), Pritchard (2015, 2018), and Moss (2021).

a specification of Thomson's (1986) causal account of individual evidence.

2 Factive Sensitivity

As pointed out in the Introduction, the idea that sensitivity marks off individual from merely statistical evidence is not new. Enoch et al. (2012) and Enoch and Fisher (2015) use the idea to analyse individual evidence. They draw upon Nozick's (1981) condition of sensitivity for knowledge. Nozick defends that knowledge is justified *true* belief that is also sensitive. He analyses sensitive belief in the context of knowledge assuming that the concept of knowledge implies true belief. An agent's true belief that p is sensitive iff, had p been false, the agent would not have believed that p . Nozick's idea is that you know p if you believe p , p is true and justified, and had p been false, you would not have believed p . Note that the subjunctive conditional asks us what you would have believed had p been false, even though p is assumed to be true and assumed to be believed.

Enoch et al. transfer Nozick's analysis of sensitive true belief to an analysis of sensitive evidence. An agent's evidence e is sensitive to a proposition p iff, had p been false, the agent would not have believed e . If e is sensitive to p , e is individual evidence for p ; otherwise e is merely statistical evidence. This is a simple and intuitive analysis of individual evidence. However, we think it faces two major problems.

First, note that Enoch et al.'s sensitivity analysis of evidence goes beyond the realm of an agent's beliefs. The antecedent 'had p been false' is a statement about facts, not beliefs. A fact finder has, of course, no access to what is actually true and false, and so no access to true beliefs or even knowledge. The actual truth or falsehood of p , as opposed to merely justified belief that p , cannot have any bearing on what verdict the court should reach. If a fact finder had access to what is actually true, there would be no problem in convicting the actually guilty and acquitting the actually innocent. Alas, a fact finder is restricted to its fallible beliefs. So the best she can do is to base her judgment only on beliefs which are more or less justified by the admissible and available evidence. And yet, whether or

not evidence is sensitive, is on Enoch et al.'s account determined by metaphysical facts the fact finder has no privileged epistemic access to.

As a result of going beyond the realm of beliefs, Enoch et al.'s notion of sensitivity does imply factivity. Their analysis of sensitive evidence inherits from Nozick's analysis of true sensitive belief that p must be true: e is sensitive to p only if p is true. To see this, assume an epistemic agent falsely believes that p . But then, if p were false (as it actually is), the epistemic agent would still believe what she actually believes, including p and any evidence e . So no false belief can be sensitive on their account.

To know that p seems to imply that p is true. But this factivity of p is at odds with a notion of evidence. Unlike knowledge, the notions of merely justified belief and evidence are not factive. You may believe p is true while it is in fact false, and you may believe p is false while it is in fact true. The evidence available to you can be misleading, and is thus likewise not factive. Hence, sensitivity – as a property of belief or evidence – should not imply factivity.

Sensitivity, as an analysis of individual evidence, is problematic in virtue of its factivity. For illustration, consider a legal case where the defendant is innocent and yet there is compelling individual but misleading evidence that the defendant is guilty. If the defendant were in fact innocent, would the court – based on the misleading evidence – believe she is innocent? Since the defendant is actually innocent and the fact finder believes she is guilty, the answer is no. The misleading evidence is not sensitive. So no wrongful conviction based on sensitive evidence is possible on Enoch et al.'s account.

On accounts of sensitivity, a defendant should be convicted only if the fact finder's judgment is based on sensitive evidence. Since only true judgments can be factually sensitive, no wrongful conviction is conceptually possible. The impossibility of wrongful convictions is an absurd consequence of the factive sensitivity account. As unfortunate as it is, wrongful convictions based on individual evidence exist. But the factive sensitivity account cannot capture such cases.

For authors like Blome-Tillmann (2015, p. 108) and Smith (2018, p. 1204), the problem of wrongful convictions is reason enough to reject the sensi-

tivity approach wholesale. We agree that a factive notion of sensitivity is untenable. However, we think the sensitivity approach can be rescued by spelling it out in purely epistemic terms, as we will do in the next section.

Here is the second problem. On Enoch et al.'s analysis, it is not clear whether individual evidence comes out as sensitive. In *Witness*, the eye-witness testimony that John climbed the fence is sensitive to John's gate-crashing iff, had John not gatecrashed, the fact finder would not have believed that the eye-witness testifies that John climbed the fence. The question is whether or not this subjunctive conditional is true.

Enoch et al. employ the standard possible worlds semantics for conditionals. A conditional $p > q$ is true at a possible world w just in case the possible worlds, which make p true and are otherwise most similar to w , are possible worlds where q is true as well (Stalnaker, 1968; Lewis, 1973b,c). A possible world is simply a possibility, a way the world might be.⁴ On this picture, a proposition p is identified with the set of those possible worlds which satisfy p . For simplicity, let us call a world where p is true a p -world and a world where p is false a $\neg p$ -world. So $p > q$ is true at a world w iff all the p -worlds most similar to w are q -worlds.

Now, consider the conditional 'if John had not gatecrashed ($\neg p$), the fact finder would not have believed the eye-witness testimony ($\neg e$).' This conditional $\neg p > \neg e$ comes out true on the standard Lewis-Stalnaker semantics just in case all the most similar $\neg p$ -worlds are $\neg e$ -worlds. For the eye-witness testimony to be sensitive, the similarity order between worlds must be specified in a way such that there is no most similar $\neg p$ -world that happens to be an e -world. But it is far from obvious how to do this.

Recall that the eye-witness is stipulated to be not fully reliable. So this is a possibility: John did not gatecrash, but the eye-witness mistook John for a gatecrasher, and testifies that John climbed the fence; and so the fact finder falsely believes that John gatecrashed. Sensitivity accounts need to specify the similarity order between worlds such that this $\neg p \wedge e$ -world is not among the most similar $\neg p$ -worlds. Otherwise the conditional $\neg p > \neg e$ comes out false and the eye-witness evidence is not sensitive.⁵

⁴For details, see Stalnaker (2003).

⁵We would like to thank an anonymous referee for bringing this possibility to our

Enoch et al. do not tell us how the similarity order over possible worlds should be specified. To be fair, they gesture at a specification by qualifying their notion of sensitivity by a ‘most probably’ clause: an agent’s evidence e is sensitive to a proposition p iff, had p been false, the agent would *most probably* not have believed e .⁶ However, it remains unclear what the ‘most probably’ clause exactly means and how it might help to specify the similarity order. To be clear, it is not obvious how the clause helps to determine the most similar $\neg p$ -worlds. And so it remains unclear how the similarity order should be specified.

The problem of how to specify a suitable similarity order between possible worlds is serious. It is, perhaps, even more difficult when similarity is not defined in terms of probability. For at least some fans of sensitivity, the underlying similarity is supposed to be of a kind ‘distinct’ from probability and cannot be reduced to it.⁷

The problem is quite general. It also haunts other accounts of individual evidence which are sometimes regarded to be superior to sensitivity analyses. Smith (2018, p. 1208), for example, analyses individual evidence in terms of *normic support*. Evidence e normically supports p iff the circumstances in which e is true and p is false would be less normal, in the sense of requiring more explanation, than the circumstance in which e and p are both true. The parallel to sensitivity analyses is striking. The notion of normic support likewise presupposes that worlds can be ranked according to their ‘comparative normalcy’. e normically supports p iff the most normal $\neg p \wedge e$ -worlds are all less normal than the most normal $p \wedge e$ -worlds. And comparative normalcy is supposed to come apart from probability comparisons.⁸ Hence, the question arises why the $\neg p \wedge e$ -worlds are less normal than the $p \wedge e$ -worlds? If John gatecrashed and the witness correctly testifies, we have a $p \wedge e$ -world. By contrast, if John did not gatecrash and the witness mistakenly testifies that he did, we have a $\neg p \wedge e$ -world. But why would the latter possible world require more ex-

attention.

⁶See Enoch et al. (2012, p. 204) and Enoch and Fisher (2015, p. 574).

⁷See, for example, Pardo (2018, fn. 10).

⁸For details on the notion of normic support, normalcy, and explanation, see Smith (2010, pp. 13-9).

planation than the former? One must not forget that the eye-witness is only partially reliable.⁹

A different approach seeks to analyse individual evidence in terms of *safe belief* (Pritchard, 2018; Pardo, 2018). Roughly, an agent has a safe belief in p iff there is no nearby world where the agent believes p but p is, in fact, false. The *Witness* scenario, however, explicitly allows for worlds where the eye-witness is mistaken and the fact finder thus falsely believes that John gatecrashed based on the misleading evidence. But then – in the absence of a reason why these worlds are not ‘nearby’ – it seems as if the belief in John’s gatecrashing is not safe. And so the eye-witness evidence does not count as individual.

In an attempt to clarify his safety account, Pritchard (2018, p. 117) writes: ‘what is required for a conviction is evidence such that, given that evidence, it cannot be an easy possibility that the defendant is wrongfully convicted.’ Given her partial reliability, it is more than a live possibility that the eye-witness is mistaken. But why is a mistaken witness not an *easy* possibility? The fans of safety accounts owe us an answer to this question.

We have seen that sensitivity analyses and their competitors face a serious problem. They need to specify criteria to determine the underlying order of possible worlds in such a way that partially reliable eye-witness evidence comes out as individual evidence. In the absence of such criteria, analyses of individual evidence in terms of sensitivity, normic support, and safe belief are at best incomplete.

3 Epistemic Sensitivity

In this section, we put forth an analysis of epistemic sensitivity. We model an agent’s belief state, relative to which we spell out an order of possible worlds. Our analysis gives rise to a schema of legal proof, which we apply to *Witness* and *Paid per Person*. Moreover, our analysis suggests a general distinction between individual and statistical evidence. Finally, we will

⁹For a criticism of normic support along similar lines, see Steele (ms.).

point out why our analysis is not susceptible to the problem of wrongful conviction.

Here is our analysis of epistemic sensitivity. An agent believes e is sensitive to p iff the agent believes two conditionals, $p > e$ and $\neg p > \neg e$. We understand the conditionals in a suppositional way. For a piece of evidence e being epistemically sensitive to a proposition p , the agent must come to believe e upon supposing p , and the agent must come to believe $\neg e$ upon supposing $\neg p$.¹⁰ Note that our analysis of sensitive evidence is purely epistemic, unlike Enoch et al.'s.

Our analysis requires a model of an agent's belief state. This model must have the resources to represent suppositions. Some suppositions have consequences that contravene an agent's non-suppositional beliefs. I believe that it does not rain and that I did not take an umbrella with me this morning. However, I also believe that I would have taken an umbrella with me this morning if it had rained. For such cases, we need more than a simple set of possible worlds to represent an agent's suppositional beliefs; in addition, we need (a set of) possible worlds the agent excludes to be actual and yet they are relevant possibilities under the supposition. In the example, the agent's belief state needs to include the possible world where it rains and I take an umbrella, even though I believe that it does not rain and I didn't take an umbrella with me.

We model an agent's belief state by a finite set $\$$ of nested sets of possible worlds and a probability distribution P that assigns non-zero probability only to the smallest element S in $\$$.¹¹ For our purposes, you may conceive of $\$$ as a set that contains two sets, S and S' , such that $S \subset S'$. The set S contains the possible worlds the agent cannot exclude to be actual. The probability distribution assigns each candidate for the actual world, or equivalently epistemic possibility, a non-zero probability value. An agent has *full belief* only in the propositions which are true in each candidate for

¹⁰Belief in the first conditional $p > e$ can be seen as an epistemic version of a safety condition: the agent believes that in all 'nearby' worlds, where p is true, the evidence e obtains. Our analysis can thus be conceived of as combining a safety and a sensitivity condition.

¹¹Formally, $\$$ is a set that contains a smallest element S and finitely many supersets S' of S . For details, see Grove (1988, p. 159).

the actual world.¹² The probability distribution P represents the agent's (non-suppositional) degrees of belief, or *credences*. The set S' contains the possible worlds that contravene the agent's current beliefs but may become relevant under a supposition.

We represent suppositions as follows. Supposing a proposition p in a belief state $\langle \$, P \rangle$ is represented by the set S^p that contains any p -world that occurs in some member of $\$$. Furthermore, supposing p results in a suppositional probability distribution P^p which is obtained from P by shifting the probability shares from the worlds in S to the worlds in S^p .¹³

Now, what does it mean that an agent believes a conditional? Well, an agent believes the conditional $p > e$ iff $p > e$ is true in her belief state $\langle \$, P \rangle$. It remains to say when a conditional is true in a belief state. We define that $p > e$ is true in $\langle \$, P \rangle$ iff all the *most likely* worlds in S^p are e -worlds. The possible worlds are simply ordered according to their suppositional probability value: the 'most similar' worlds are the most likely worlds under the supposition. Our semantics of conditionals relies on a notion of 'comparative probability' – if you like. The idea is that only the most likely hypothetical scenarios, where the antecedent is true, should be considered to evaluate the consequent. In sum, an agent believes $p > e$ iff upon supposing all the most likely scenarios where p is true the agent would come to believe that e is true.

We think of evidence as any kind of information that structures what an agent considers possible and impacts her beliefs and credences. In light of new evidence, you may obtain new beliefs, give up previously held beliefs, or change your credences in certain propositions. Importantly, evidence also shapes how we conceive of the structure of a given scenario: which possibilities we discern, and how we distribute probabilities over those possibilities. As a methodological rule, we will use as many possi-

¹²On our model, an agent neither believes p nor $\neg p$ iff p is true at some of the worlds in S and $\neg p$ is true at some other worlds in S . As a consequence, our agents believe all logical truths and perhaps all necessary truths. This is, of course, an idealisation. But we think this idealisation is not overly harmful. Nobody will be convicted because a fact finder believes a tautology.

¹³Technically speaking, P^p is the *general image* of P on p . For details, see Lewis (1976, pp. 308-12) and Gärdenfors (1982, p. 751).

ble worlds as necessary to represent the evidence an agent obtained, but no more. We will represent the possibilities made relevant by the evidence by a set of possible worlds W that is as small as possible without losing any information. Quite naturally, the set W of relevant possibilities will be the largest element of \mathcal{S} in the respective belief state. On this picture, the obtained evidence determines the space of epistemic possibilities, including both candidates for the actual world and mere hypothetical possibilities, and assigns the candidates for the actual world particular probability values. In brief, the full beliefs, credences, and epistemic possibilities of our agent are structured by the obtained evidence.

Let us return to the legal case. Our agent is a fact finder, e stands for the available and admissible evidence, and p for a legal claim. We propose thus: a fact finder believes evidence e is sensitive to a claim p iff the agent believes that the evidence e would obtain if the claim p were true, and the evidence e would not obtain if the claim p were false. The conditionals are understood as above: upon supposing the most likely p -scenarios, the fact finder would come to believe e , and upon supposing the most likely $\neg p$ -scenarios, the fact finder would come to believe $\neg e$.

Our analysis of epistemic sensitivity gives rise to a schema of legal proof. Where P represents a fact finder's credences, s a threshold value, p a claim, and e the admissible and available evidence, a defendant should be found reliable for p iff $P(p | e) > s$ and the fact finder believes e is sensitive to p . In civil cases, s is $1/2$. In criminal cases, s can usually be determined by a classic decision-theoretic argument.¹⁴

3.1 *Witness and Paid Per Person Revisited*

Recall the *Witness* scenario. The key evidence is a fairly reliable eyewitness who testifies that John climbed the fence. It is uncontested that the testimony makes it 70% probable that John gatecrashed. This evidence structures our fact finder's beliefs, credences, and epistemic possibilities.

¹⁴The decision-theoretic argument fixes the numerical threshold s such that convicting has higher expected value than acquitting. For details, see for example, Cheng (2013, pp. 1259-61&1275-8), Steele (ms., pp. 3-4), and Di Bello (2019, pp. 1054-5).

The fact finder believes that the eye-witness testified that John climbed the fence. She does neither believe that John gatecrashed nor that he did not. After all, the eye-witness is not perfectly reliable and so she thinks it is possible that the witness testifies that John climbed the fence when, in fact, John did not gatecrash. Perhaps another man climbed the fence and the witness mistook this man for John. Still, the fact finder has a credence of 70% that John gatecrashed.

Crucially, an agent's reliability goes over and above actuality. A perfectly reliable agent reports p in a hypothetical scenario where p is the case, and reports $\neg p$ in a hypothetical scenario where $\neg p$ is the case. The less reliable an agent, the less often the agent's report matches what is the case in particular (hypothetical) scenarios. In virtue of this modal character, the reliability of the eye-witness shapes even the fact finder's belief-contravening possibilities. The fact finder believes that the eye-witness testified that John climbed the fence. However, the reliability carries information about possibilities which conflict with what the fact finder considers to be actual. The witness's reliability determines that the fact finder has the following hypothetical belief: it is more likely that the witness does not testify that John climbed the fence when he did not gatecrash than when he did gatecrash. So the eye-witness evidence determines that the fact finder has merely hypothetical opinions on possibilities where the witness does not testify.

Should John be found liable for gatecrashing? According to our account of legal proof, yes. To see this, let p stand for 'John gatecrashed' and e for 'the eye-witness testified that John climbed the fence'. Our fact finder's $\$$ can be represented by the following box:

$w_1 : p \wedge e$	$w_2 : \neg p \wedge e$
$w_3 : p \wedge \neg e$	$w_4 : \neg p \wedge \neg e$

The upper layer represents $S = \{w_1, w_2\}$, and the lower and upper layer together $S' = \{w_1, w_2, w_3, w_4\}$.¹⁵ The witness's good but imperfect reliability is a modal notion. Across (merely hypothetical) scenarios, her reliability makes it more likely that her reports say what is the case. Hence,

¹⁵ $\$ = \{S, S'\}$.

our fact finder believes the ‘matching’ worlds, where the truth value of p and e match, to be more likely than the ‘non-matching’ worlds, where the truth values do not match.

Our fact finder believes e is sensitive to p iff $p > e$ and $\neg p > \neg e$ are both true in her belief state $\langle \$, P \rangle$. $p > e$ is true in $\langle \$, P \rangle$ iff all the most likely worlds in S^p are e -worlds. The set S^p contains the worlds w_1 and w_3 . w_1 is a matching world and w_3 is not. So w_1 is more likely than w_3 according to the suppositional distribution P^p .¹⁶ Hence, $p > e$ is true in $\langle \$, P \rangle$. Our analysis captures the intuition that the fact finder believes: if John gatecrashed, the witness would (in the most likely scenario) have testified that John climbed the fence.

Similarly, $\neg p > \neg e$ is true in $\langle \$, P \rangle$ iff all the most likely worlds in $S^{\neg p}$ are $\neg e$ -worlds. The set $S^{\neg p}$ contains the worlds w_2 and w_4 . w_4 is a matching world and w_2 is not. So w_4 is more likely than w_2 according to the suppositional distribution $P^{\neg p}$. Hence, $\neg p > \neg e$ is true in $\langle \$, P \rangle$. Our analysis captures the intuition that the fact finder believes: if John did not gatecrash, the witness would (in the most likely scenario) not have testified that John climbed the fence. We have established this: on our analysis, the fact finder indeed believes that the eye-witness testimony is sensitive to John’s gatecrashing.

According to our account of legal proof, John should be found liable for p iff $P(p | e) > 1/2$ and the fact finder believes e is sensitive to p . The conditional credence $P(p | e)$ that John gatecrashed given that the eye-witness testified that he climbed the fence is $7/10$ and so exceeds the threshold for finding liable. And since the fact finder believes the eye-witness testimony to be sensitive to John’s gatecrashing, our account recommends that John should be found liable for gatecrashing.

Should John be found liable in *Paid per Person*? This time our account of legal proof says no. Of course, the credence condition is met. The proba-

¹⁶Initially, w_1 is assigned a higher probability value than w_2 . That is, $P(p | e) = P(w_1) = 7/10$ and so, by the probability calculus, $P(\neg p | e) = P(w_2) = 3/10$. By supposing p , we consider the general image on p . Since the eye-witness is more reliable than unreliable, $P^p(w_1) > P^p(w_3)$. Hence, the unique most likely p -world in S^p according to P^p is w_1 .

bility that John gatecrashed is 70%, just like in *Witness*. However, we will show that, on our analysis, the statistical evidence is not sensitive to John’s gatecrashing.

The key evidence in *Paid per Person* is that 700 out of 1000 attendees of the local rodeo gatecrashed. There are $\binom{1000}{700}$ ways how 700 out of 1000 attendees gatecrashed.¹⁷ Since we have no more evidence, all of those very many possibilities are equally likely. The 700:300 statistics structures our fact finder’s beliefs, credences, and epistemic possibilities accordingly. The fact finder thus considers very many possible worlds w_j ($0 \leq j \leq \binom{1000}{700}$), where attendee i ($0 \leq i \leq 1000$) together with 299 select attendees did not gatecrash and all other attendees gatecrashed. And each of those worlds is assigned a (very small) positive and equal probability value. The statistics informs the fact finder that there are very many equally likely possibilities and in 70% of them a randomly picked attendee gatecrashed, while in 30% of them the attendee did not. The fact finder thus believes the statistics and has a credence of 70% that a randomly picked attendee – like John – gatecrashed.

The fact finder does not believe that the 700:300 statistics is sensitive to John’s gatecrashing. To see this, let p stand for ‘John gatecrashed’ and $stat[700 : 300]$ for ‘700 out of 1000 attendees gatecrashed’. Our fact finder’s $\$$ can be represented by the following box:

$$\boxed{w_1, w_2, \dots, w_j}$$

The layer represents S .¹⁸ $stat[700 : 300]$ is true in all the worlds in S , and in 70% of those worlds p is true, in 30% $\neg p$. Unlike the partially reliable witness testimony, the 700:300 statistics carries no modal information. The witness reliability specifies how likely the witness report would be if John did or did not gatecrash, respectively. By contrast, the statistics on its own does not specify how likely it would be if John did or did not gatecrash. The statistics is fully believed and it is silent on belief-contravening possibilities.

¹⁷The binomial coefficient $\binom{1000}{700}$ equals approximately $5.29 \cdot 10^{263}$.

¹⁸ $\$ = \{S\}$.

It is tempting to think that the 700:300 statistics is modal in an obvious way. Our fact finder putatively believes: had John not gatecrashed, the statistics would have been 699:301.¹⁹ However, there are two problems with this thought. First, it assumes that the fact finder fully believes that John actually gatecrashed – which she doesn't. She only believes that John may or may not be one of the gatecrashers. When she considers the second possibility where John actually did not gatecrash, she believes: had John not gatecrashed, the statistics would have been 700:300. Our fact finder, therefore, does not believe the above conditional: upon supposing a most likely scenario that John did not gatecrash, she still believes the 700:300 statistics.

The second problem with the tempting thought is simply this: we have no modal information. So, we do not know how the statistics would change if John had not gatecrashed. Perhaps, John was the first gatecrasher inciting the mass gatecrashing. If this information were available to our fact finder, she would not believe that the statistics changes to 699:301 if John had not gatecrashed – she would rather believe that nobody or at least much less people would have gatecrashed. And so the statistics, in conjunction with the information that John incited the mass gatecrashing, would become sensitive evidence. In the absence of such modal information, however, it is underdetermined how the statistics would have changed.

To be explicit, our fact finder believes $stat[700 : 300]$ is sensitive to p iff $p > stat[700 : 300]$ and $\neg p > \neg stat[700 : 300]$ are both true in her belief state $\langle \$, P \rangle$. The set S^p contains all the equally likely worlds where John belongs to the 700 gatecrashers. In all of them, $stat[700 : 300]$ is true. Similarly, the set $S^{\neg p}$ contains all equally likely worlds where John does not belong to the gatecrashers. In those worlds, however, $stat[700 : 300]$ is still true. There are 'most likely' worlds that satisfy $\neg p$ and $stat[700 : 300]$. Hence, $\neg p > \neg stat[700 : 300]$ is not true in $\langle \$, P \rangle$, and so the statistics is not sensitive to p .

Our analysis of epistemic sensitivity explains Pritchard's intuition that John could *easily* be among the 30% who paid for admission. On our account, an easy possibility is simply one an epistemic agent considers to

¹⁹See Blome-Tillmann (2015, p. 105), Di Bello (2019, p. 1050).

be among the most likely possibilities. And indeed, the p - and $\neg p$ -worlds are equally likely according to the statistics, and so there are most likely $\neg p$ -worlds. This is, on our analysis, the reason why it is salient to believe that John could be among the non-gatecrashers. Although it is less likely that John paid for admission, this possibility could happen just as well as any other.

The gatecrasher scenarios suggest a general distinction between statistical and individual evidence. Statistical evidence contains information about usually many possibilities and a uniform (or near uniform) probability distribution over those possibilities. This induces a certain symmetry between the possibilities: each possibility is just as likely as any other. Statistical evidence does not probabilistically discriminate between this or that particular possibility. The 700:300 statistics, for example, says nothing special about John. Anything it says applies equally to the other attendees. In other words, all the statistics says about John it says likewise about any other attendee in the stadium. In this sense, statistical evidence does not discriminate between its possibilities.

Individual evidence, by contrast, contains information only about the particular case at hand. The partially reliable eye-witness testimony is only about John. It does not tell us anything on the other attendees, let alone whether or not they gatecrashed. The individual evidence is silent on the other attendees. However, the testimony has a modal force on the fact finder's beliefs regarding the specific case of John: it renders relevant the four possibilities w_1 - w_4 . And more importantly, the eye-witness testimony allows to discriminate between (hypothetical) possibilities of different (suppositional) probability. Individual evidence discriminates between its relevant possibilities.²⁰

²⁰It follows that eye-witness testimony need not be individual according to our analysis of epistemic sensitivity. Suppose a witness is a perfect randomizer: she reports p in any hypothetical scenario, where p is the case, in exactly half of the cases, and she reports $\neg p$ in any hypothetical scenario, where $\neg p$ is the case, in exactly half of the cases. Then her testimony is epistemically insensitive on our analysis. And indeed, her testimony is random and cannot be used – not even for a 'stop-and-frisk' search.

3.2 Wrongful Convictions are Possible

Finally, let us explain why our analysis of epistemic sensitivity allows for wrongful convictions. The underlying reason is that our analysis is purely epistemic, it does not go beyond the realm of beliefs. Hence, epistemic sensitivity does not imply factivity: a piece of evidence e may be epistemically sensitive to p even if p is false. To see this, take a case where the claim p , viz. the defendant is guilty, is in fact false. However, there is plenty misleading evidence e that raises our fact finder's credence $P(p)$ over the threshold s required for conviction. Based on the compelling and misleading evidence e for p , our fact finder would falsely convict if e is epistemically sensitive to p . And indeed, it is still a conceptual possibility that the fact finder falsely believes $p > e$ and $\neg p > \neg e$. It is an open question whether the fact finder would believe e upon supposing the most likely p -scenarios. It is likewise an open question whether the fact finder would believe $\neg e$ upon supposing the most likely $\neg p$ -scenarios. So, even though p is false in this case, there can be epistemically sensitive evidence e in support of p .

Our analysis of epistemic sensitivity is thus not susceptible to the problem that wrongful convictions are conceptually impossible. Indeed, innocent defendants should be convicted on our account just in case the fact finder's credence in the defendant's guilt given the misleading evidence exceeds the appropriate threshold and the fact finder believes that the misleading evidence is sensitive to the claim that the defendant is guilty. Since epistemic sensitivity does not go beyond the realm of beliefs, wrongful convictions are a conceptual possibility on our account, as desired.

4 Would-Be Causation

Our analysis of epistemic sensitivity is related to Thomson's (1986) account of individual evidence: our analysis can be seen as a specification of her account. She proposed to analyse individual evidence in terms of causation. She says, roughly, that an agent's evidence e in support of a proposition p is individual iff the agent believes that p is a would-be cause

of e .²¹

The idea that causation is the mark of individual evidence is intuitively appealing. We believe that John's gatecrashing would be a cause of the eye-witness seeing John climb the fence. By contrast, we do not believe that John's gatecrashing would cause that 70% of attendees gatecrash (unless we believe that John was the first gatecrasher causing the mass gatecrashing). In general, we believe that individual evidence is caused by a particular event or fact, but we do not believe that a particular event or fact causes statistical evidence. This is how the causal account of individual evidence solves the problem of statistical evidence on an intuitive level.

Thomson's account has been often and severely criticised because she does not spell out the causal relation that figures so prominently in her account (see Redmayne (2008), Enoch et al. (2012), Enoch and Fisher (2015), Gardiner (2018)). Now, if we specify Thomson's cause as a Lewisian difference maker understood in terms of our conditionals, our analysis of epistemic sensitivity coincides with her causal account. Or so we will argue in the remainder of this section.

Lewis (1973a) thinks of causes as difference makers. Whether or not a cause occurs makes a difference as to whether or not its effect occurs. He thus proposed to analyse causation between particular events in terms of counterfactual conditionals. If p and e occur, and the counterfactual if p had not occurred, e would not have occurred is true, his analysis implies that p is a cause of e . Together with the factivity of p and e , the counterfactual expresses that p makes a difference as to whether or not e occurs.

Unlike Lewis, who is after a metaphysical relation of causation, we should not presuppose in Thomson's causal relation that p is true, or that the fact finder under consideration already believes p . As we have seen above, the factivity of a claim p leads to the problem that wrongful convictions are impossible; and if the fact finder is already convinced that p , the legal

²¹Here, we go over concerns of forward- and backward-looking individual evidence. Moreover, I presume that Thomson's account is purely epistemic. The textual evidence is not definite. If her account is factive, wrongful convictions are conceptually impossible, just like for Enoch et al.'s sensitivity account.

proceedings are superfluous. Our question is rather whether we believe that p could be a cause that would make a difference as to the evidence we currently believe. If so, our evidence is individual relative to p ; if not, not.

Now, p is a would-be difference maker as to e iff two subjunctive conditionals are true: $p > e$ and $\neg p > \neg e$. Hence, an agent believes p is a would-be difference maker of e iff the agent believes $p > e$ and $\neg p > \neg e$. Assuming our semantics of conditionals and that Thomson's causes are Lewisian difference makers, we obtain the following result: e is epistemically sensitive to p iff e is individual evidence for p on Thomson's causal account.

Of course, we do not know whether Thomson had Lewisian difference makers in mind. But if we understand her causes as difference makers in terms of our conditionals, we end up with our analysis of epistemic sensitivity for which we have shown that it solves the gatecrasher scenarios and allows for wrongful convictions.

Spelling out causes as difference makers sheds light on what Thomson (1986, pp. 212-15) might mean by 'guarantee'. She says a causal relation provides us with a guarantee that our belief in a legal claim is not merely true as a matter of luck. If a fact finder decides that John is guilty of gatecrashing merely based on the statistical evidence that 70% of attendees gatecrashed, then it is just luck if she happens to justly impose liability on John. It is still an easy (or most likely) possibility for her that John did not gatecrash. By contrast, if she decides that John is guilty of gatecrashing based on the evidence that an eye-witness saw John climb the fence, she has a justification of why the imposition of liability is just. She believes that whether or not John gatecrashed makes a difference as to the eye-witness testimony. And so, upon receiving the eye-witness testimony, it is not an easy (or most likely) possibility anymore for her that John did not gatecrash. The eye-witness evidence guarantees John's gatecrashing in this sense.

In general, a piece of evidence e can be understood as a guarantee for a legal claim p and so as making e individual iff p is believed to be a difference maker for e in terms of our conditionals. To require a guarantee for a conviction can thus be understood as follows: the fact finder believes that

the available evidence is most likely if the defendant committed the crime, and the evidence does not occur in the most likely scenarios where the defendant did not commit the crime. This requirement expresses a justification for convictions different from a merely high credence. Hence, our analysis of epistemic sensitivity – just like Thomson’s account based on difference-making causation – prevents that a defendant is convicted on ‘purely statistical grounds’ without any justification of another kind. The requirement helps avoid, for instance, that a defendant is convicted only because the defendant belongs to a demographic group which is disproportionately represented in certain crimes.

5 Conclusion

We have put forth an analysis of epistemic sensitivity. Essentially, an agent believes e is sensitive to p iff the agent believes two conditionals, $p > e$ and $\neg p > \neg e$. We have detailed when an agent believes a conditional. In particular, we have modelled an agent’s belief state relative to which we have specified an order of possible worlds in terms of their suppositional probability values. An agent believes $p > e$ iff upon supposing all the most likely scenarios where p is true the agent would come to believe that e is true.

Our analysis of epistemic sensitivity gives rise to an account of legal proof. Where P represents a fact finder’s credences, p a legal claim, and e the admissible and available evidence, a defendant should be found liable for p iff $P(p \mid e) > 1/2$ and the fact finder believes e is sensitive to p . Although the fact finder’s credence that John gatecrashed has the same value in *Witness* and *Paid per Person*, this account of legal proof finds liable in the former but not in the latter scenario. For the eye-witness evidence is epistemically sensitive on our analysis, whereas the 700:300 statistics is not. Moreover, unlike the account of Enoch et al. based on factive sensitivity, our account allows for the conceptual possibility of wrongful convictions.

Finally, we have outlined how our analysis of epistemic sensitivity can be understood as a causal account of individual evidence. In particular, our analysis can be seen as a specification of Thomson’s causal account if

causes are conceived of as Lewisian difference makers defined in terms of our conditional. We have seen that our analysis of epistemic sensitivity and its causal clothing suggest a distinction between statistical and individual evidence. However, in order to draw a neat distinction, we still need to consider mixed bodies of evidence that contain individual alongside statistical pieces of information. This investigation has to wait for another occasion.

Acknowledgements. This paper would not exist were it not for Katie Steele who introduced me to its topic and provided insightful and extensive comments. I am grateful for detailed comments by Atoosa Kasirzadeh, Alan Hájek, Pamela Robinson, and the team of editors and referees of *Inquiry*. Furthermore, I would like to thank the audience at a Humanising Machine Intelligence seminar for their valuable feedback.

References

- Blome-Tillmann, M. (2015). Sensitivity, Causality, and Statistical Evidence in Courts of Law. *Thought: A Journal of Philosophy* 4(2): 102–112.
- Cheng, E. (2013). Reconceptualizing the Burden of Proof. *Yale Law Journal* 122(5): 1254–1279.
- Cohen, L. J. (1977). *The Probable and the Provable*. Oxford: Clarendon Press.
- Di Bello, M. (2019). Trial by Statistics: Is a High Probability of Guilt Enough to Convict? *Mind* 128(512): 1045–1084.
- Enoch, D. and Fisher, T. (2015). Sense and ‘Sensitivity’: Epistemic and Instrumental Approaches to Statistical Evidence. *Stanford Law Review* 67: 557–611.
- Enoch, D., Spectre, L., and Fisher, T. (2012). Statistical Evidence, Sensitivity, and the Legal Value of Knowledge. *Philosophy & Public Affairs* 40(3): 197–224.
- Gärdenfors, P. (1982). Imaging and Conditionalization. *The Journal of Philosophy* 79(12): 747–760.

- Gardiner, G. (2018). Legal Burdens of Proof and Statistical Evidence. In *The Routledge Handbook of Applied Epistemology*, edited by J. Chase and D. Coady, Routledge.
- Grove, A. (1988). Two modellings for theory change. *Journal of Philosophical Logic* **17**: 157–170.
- Haack, S. (2014). *Evidence Matters: Science, Proof, and Truth in the Law*. Law in Context, Cambridge University Press.
- Lewis, D. (1973a). Causation. *Journal of Philosophy* **70**(17): 556–567.
- Lewis, D. (1973b). *Counterfactuals*. Oxford: Blackwell.
- Lewis, D. (1973c). Counterfactuals and Comparative Possibility. *Journal of Philosophical Logic* : pp. 418–446.
- Lewis, D. (1976). Probabilities of Conditionals and Conditional Probabilities. *The Philosophical Review* **85**(3): 297–315.
- Moss, S. (2021). Knowledge and Legal Proof. In *Oxford Studies in Epistemology*, Oxford: Oxford University Press.
- Nozick, R. (1981). *Philosophical Explanations*. Oxford: Oxford University Press.
- Pardo, M. S. (2018). Safety vs. Sensitivity: Possible Worlds and the Law of Evidence. *Legal Theory* **24**(1): 50–75.
- Pritchard, D. (2015). Risk. *Metaphilosophy* **46**(3): 436–461.
- Pritchard, D. (2018). Legal risk, legal evidence and the arithmetic of criminal justice. *Jurisprudence* **9**(1): 108–119.
- Redmayne, M. (2008). Exploring the Proof Paradoxes. *Legal Theory* **14**(4): 281–309.
- Ross, L. (2020). Rehabilitating Statistical Evidence. *Philosophy and Phenomenological Research* doi:10.1111/phpr.12622.
- Smith, M. (2010). What Else Justification Could Be. *Noûs* **44**(1): 10–31.

- Smith, M. (2018). When Does Evidence Suffice for Conviction? *Mind* 127(508): 1193–1218.
- Smith, M. (2020). Against Legal Probabilism. In *Truth and Trial*, edited by J. Robson and Z. Hoskins, Routledge.
- Stalnaker, R. (1968). A Theory of Conditionals. In *Studies in Logical Theory (American Philosophical Quarterly Monograph Series)*, edited by N. Rescher, no. 2, Oxford: Blackwell. 98–112.
- Stalnaker, R. (2003). *Ways a World Might Be*. Oxford: Clarendon Press.
- Steele, K. (ms.). Evidence, Arbitrariness, and Fair Treatment. *Unpublished Manuscript* : 1–21.
- Thomson, J. J. (1986). Liability and Individualized Evidence. *Law and Contemporary Problems* 49(3): 199–219.