

# Knowledge, Individualised Evidence and Luck

Dario Mortini

(Forthcoming in *Philosophical Studies*)

## Abstract

The notion of individualised evidence holds the key to solve the puzzle of statistical evidence, but there's still no consensus on how exactly to define it. To make progress on the problem, epistemologists have proposed various accounts of individualised evidence in terms of causal or modal anti-luck conditions on knowledge like appropriate causation (Thomson 1986), sensitivity (Enoch et al. 2012) and safety (Pritchard 2018). In this paper, I show that each of these fails as satisfactory anti-luck condition, and that such failure lends abductive support to the following conclusion: once the familiar anti-luck intuition on knowledge is extended to individualised evidence, no single causal or modal anti-luck condition on knowledge can succeed as the right anti-luck condition on individualised evidence. This conclusion casts serious doubts on the fruitfulness of the move from anti-luck conditions on knowledge to anti-luck conditions on individualised evidence. I expand on these doubts and point out further aspects where epistemology and the law come apart: epistemic anti-luck conditions on knowledge do not adequately characterise the legal notion of individualised evidence.

**Keywords:** Legal Epistemology - Individualised Evidence - Epistemic Luck

## Introduction

A witness, a tape, a video, a trace: this is evidence *individualised* to a fact. It's not very reliable: witnesses lie or err, tapes can be faked, images are ambiguous, and traces mislead. Statistics, percentages, estimations of probability: this is *bare* statistical evidence. It seems more reliable: it can make crucial incriminating facts highly probable, and seldom leads astray. Yet, the law appears to prefer the former over the latter. But why? What's so special about individualised evidence? These are the main questions raised by the puzzle of statistical evidence, and while the notion of individualised evidence may hold the key to solving it, there's still no agreement on how exactly to define it. In this passage, Judith Thomson gives us a crucial hint:

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, and they flip a coin, which comes up heads. So they return and say, "the defendant is guilty." *Their so doing is not made acceptable by the fact (supposing it a fact) that the defendant actually is guilty.* (...) It matters to us, not just that a defendant does not suffer a penalty unjustly, but also that the penalty is not imposed on him unjustly. (Thomson 1986: 213, emphasis added)

Her remarks are significant, and suggestive. Justice demands that punishment and liability are imposed on the basis of the right evidence, and the right evidence is incompatible with (a certain type of) luck. It's thus an important desideratum that individualised evidence satisfies an anti-luck condition. The point is familiar to epistemologists: an identical desideratum applies to knowledge. Indeed, the main challenge underlying the Gettier problem consists in identifying the right anti-luck condition on knowledge, as Duncan Pritchard (2015: 94) points out here:

Gettier's famous paper didn't just demonstrate that knowledge wasn't (non-factively) justified true belief. He also demonstrated that one could have a justified true belief which was nonetheless subject to knowledge-undermining epistemic luck. This raises a challenge regarding what condition or conditions must be imposed on knowledge in order to exclude such luck. Call this formulation of the Gettier problem the anti-luck problem.

The anti-luck problem seems to affect both knowledge and individualised evidence, and indicates a deeper structural connection between the two. This connection did not go entirely unnoticed: epistemologists have offered accounts of individualised evidence in terms of epistemic qualities which are also advanced as anti-luck conditions on knowledge. However, when imposed

on knowledge, these conditions are unsuccessful: they face perfectly general and possibly unavoidable difficulties. In this paper, I take up the task to show that the same difficulties carry over to individualised evidence: my conclusion will be that, given the difficulties incurred with knowledge, no single causal or modal anti-luck condition on knowledge can succeed as the right anti-luck condition on individualised evidence. And this is a potentially significant result, for it casts into serious doubt the fruitfulness of the move from an anti-luck condition on knowledge to an anti-luck condition on individualised evidence.

My plan is as follows. In §1, I set the stage and lay out the puzzle in detail. In §2, I clarify further the connection between knowledge, individualised evidence and luck. Next, in §3 and §4, I offer Gettier-style variations of stylized court cases and show that prominent causal and modal anti-luck conditions on knowledge fail also when imposed on individualised evidence. Finally, in §5, I advance a diagnosis of the failure of these prominent anti-luck conditions. I expand on the relevance of such failure and I point out key aspects in which epistemology and the law come apart; this has important repercussions for legal epistemology in general.

## §1 The Puzzle of Statistical Evidence

In a fair trial, defendants are not required to prove their innocence: the onus of proof always falls on the party who first takes legal action. Such asymmetry stems from what is considered the golden thread of many contemporary criminal justice systems, the presumption of innocence. The *burden of proof* is so established: it is a necessary condition for conviction that the State proves the guilt of the defendant, who famously remains innocent “until proven guilty”.<sup>1</sup>

The burden of proof generates a *standard of proof*. While the current standard for civil liability is the “preponderance of evidence”, the standard for criminal conviction is the more demanding proof “beyond reasonable doubt”. These formulations are renowned, but their exact meaning is far from clear. Actually, it appears to be surprisingly *elusive*: explaining the precise nature of the standards is no easy task even for trained legal officials. As chief judge Jon Newman (1993: 984) effectively put it, “I find it rather unsettling that we use a formulation that we believe will become less clear the more we explain it”.

Among legal theorists, it is sometimes assumed that probability theory can provide a way around this conundrum. A once very popular view in legal scholarship favours a *quantitative*

---

<sup>1</sup>Scots law makes an exception, and permits a third verdict (“not proven”) in addition to “guilty” or “not guilty”. If this possibility obtains, then the defendant does not remain innocent until proven guilty.

interpretation: roughly, the civil and criminal standards neatly reduce to quantifiable *probabilistic thresholds*.<sup>2</sup> Following Blome-Tillman (2017), they can be defined in terms of evidential probabilities:

**Preponderance of Evidence:** Proposition  $p$  meets the standard iff  $P(p|e) > 0.5$

**Beyond Reasonable Doubt:** Proposition  $p$  meets the standard iff  $P(p|e) \geq 0.9$

So understood, the preponderance of evidence and beyond reasonable doubt standards are satisfied so long as, given the admissible evidence adduced to the case, the probability of liability and guilt is at least 0.6 or 0.9 respectively. The key motivation for this probabilistic interpretation is to ensure the accuracy of legal verdicts and secure the Blackstone ratio, according to which the percentage of false convictions should be significantly lower than false acquittals.

So far, so good. However, this probabilistic interpretation leads to a puzzle long discussed in evidence law – the puzzle of statistical evidence.<sup>3</sup> Let’s consider the following two cases, ranging over the civil and criminal standards:

ACCIDENT.<sup>4</sup> While driving back home, a vehicle suddenly hits Jane’s car. In the accident, she breaks her leg. She realises that the vehicle was a bus, but she cannot discern the colour. Jane runs a quick search, and discovers that the Blue Bus Company operates 80% of buses in town. The remaining 20% are operated by the Red Bus Company. Jane sues the Blue Bus Company. Using only the evidence described here, Jane expects to win the case. After all, given the evidence, it’s 0.8 probable that the Blue Bus Company is liable.

PRISONERS.<sup>5</sup> 100 prisoners are exercising in the prison yard. Suddenly 99 of them attack the guard, putting into action a plan that the 100th prisoner knew nothing about. The 100th prisoner played no role in the assault and could have done nothing to stop it. There is no further information that can be used to settle the question of any particular prisoner’s involvement. Despite the paucity of evidence, the guards pick prisoner 54 and convict him for the attack. After all, given the evidence, it’s 0.99 probable that prisoner 54 is guilty.

---

<sup>2</sup>Using Haack’s established terminology, this view goes under the name of legal probabilism and it has both prominent defenders (e.g., Calabresi 1961, Becker 1968, Posner 1973, Hedden and Colyvan 2019) and prominent critics (e.g., Tribe 1971, Cohen 1977, Allen and Leiter 2001, Allen and Stein 2013, Haack 2014). See Di Bello and Urbaniak (2021) for a thorough and updated discussion of legal probabilism. The “once” proviso is especially important: these days, theorists of evidence law have moved away from legal probabilism and endorsed an explanationist interpretation of legal proof in terms of relative plausibility. See Allen (2008), Pardo and Allen (2008, 2019) for an exhaustive presentation of an explanationist account of the main standards of proof.

<sup>3</sup>The puzzle is also known as the proof paradox. See Gardiner (2019; forthcoming) Ross (2020b), Pardo (2019: §3), Ho (2008: 135-142) and Redmayne (2008: §1) for an overview.

<sup>4</sup>Redmayne (2008), Enoch et al. (2012), Smith (2018), Gardiner (2019) discuss this case.

<sup>5</sup>Originally described in Nesson (1979), here I’m adopting the presentation from Redmayne (2008).

A puzzle looms ahead. While the evidence presented in statistical form meets the probabilistic standards, it also seems insufficient to undergird a positive legal verdict. In the vignettes above, imposing liability on the Blue Bus Company and convicting prisoner 54 seems unjust.<sup>6</sup> Conversely, a different type of evidence - individualised evidence - changes our reaction to the cases. The notion of individualised evidence takes centre stage in this paper, but it has proven surprisingly resistant to a precise definition (I return to this in §3). Hence, at this preliminary stage I must appeal to canonical examples to make it clearer.<sup>7</sup> Suppose that in ACCIDENT, Jane's usually reliable and trustworthy friend Jack *tells* her that he saw a blue bus hitting her car. Or imagine someone *seeing* prisoner 54 assaulting the guard. Further suppose that, after running several tests, it turns out that the eye-witnesses provide accurate identifications only 70% of the time in similar lighting conditions. On such basis, verdicts of liability and guilt are more acceptable, but neither piece of individualised evidence is as probative as the market share of the two companies, or the stupendously high probability that prisoner 54 assaulted the guard. As respectable psychological studies show, eye-witnesses are only imperfectly reliable.<sup>8</sup>

The puzzle gets closer. On the basis of canonical pieces of individualised evidence, the probability of liability and guilt is lower than it would be if only statistical evidence was adduced. Individualised evidence would even fail to meet the probabilistic threshold required for criminal conviction: given the reliability of eye-witnesses, it is only (roughly) 0.7 probable that prisoner 54 attacked the guard. Yet such evidence seems better suited than mere estimations of probability for securing a positive verdict. This squares badly with the probabilistic interpretation of the standards: why do we intuitively eschew highly reliable statistical evidence but accept imperfectly reliable individualised evidence?

This question is neither just a matter of subjective intuitions nor an instance of abstract philosophical speculation: legal practice also shows a rather conflicted attitude towards purely statistical evidence.<sup>9</sup> ACCIDENT is very similar (almost identical) to an actual civil lawsuit,

---

<sup>6</sup>The use of toy-cases like these may muddy the waters. As Steele and Colyvan (2022) emphasise, the presentation in PRISONERS is problematic because the intuitions of unfairness that it prompts track more equity issues rather than the use of statistical evidence: since each prisoner has the same probability of being guilty, it's unfair to randomly single out and punish only one prisoner. To better appreciate the point, imagine a line of cars that are all illegally parked. A parking inspector arrives, sees and knows that they are illegally parked but fines only one of the owners. This is unfair because of inequity: one car owner is unreasonably singled out (thanks to a referee for raising the case). Perhaps the PRISONERS case can be changed so that each of the 100 prisoners is punished, but then it's no longer clear that securing 99 known to be correct convictions is unfair (many would take this to be a very good outcome). These observations bring out the complexity underlying these toy-cases and anticipate the problems with using them to advance the debate on statistical evidence (more on this in §5).

<sup>7</sup>See Enoch and Fisher (2015: §1) and Pardo (2019: 262-266) for discussion of extant attempts to distinguish statistical from individualised evidence.

<sup>8</sup>Kahneman and Tversky (1982), Coady (1992: Chapter 15), Loftus (1996; 2019: §3).

<sup>9</sup>This is an empirical observation about the law, so I must proceed carefully. Legal practice is usually averse

*Smith v. Rapid Transit Inc.* Unsurprisingly, the bus company was *not* found liable. Moreover, studies involving mock juries suggest that trained judges are equally reluctant to impose liability using statistical evidence alone, and also have a preference for individualised evidence (e.g., Wells 1992).<sup>10</sup> Such antipathy is best appreciated in terms of *sufficiency* rather than *admissibility*: statistical evidence is not inadmissible, but legal officials seem hesitant to convict solely on the basis of it.<sup>11</sup>

The puzzle is here. On the one hand, according to a probabilistic interpretation, statistical evidence should be sufficient to meet the standards, but both intuitions and legal practice deny this. On the other, according to a probabilistic interpretation, individualised evidence should often be insufficient to meet the standards, but both intuitions and legal practice deny this too. We face a puzzle: the puzzle of statistical evidence.

## §2 A Common Anti-Luck Condition

The puzzle has been long discussed in evidence law, but it has recently sparked a lively debate also in epistemology.<sup>12</sup> On the negative side, legal epistemologists have rejected a probabilistic interpretation of the standards.<sup>13</sup> More positively, they have proposed to identify individualised evidence in light of specific epistemic qualities usually imposed as necessary conditions. At this point in the picture, enter knowledge: crucially, the epistemic qualities imposed as necessary

---

to *certain* pieces of statistical evidence, but not to others, such as DNA evidence. DNA evidence is presented in statistical form, and yet it has helped to secure many post-conviction DNA exonerations. See Roth (2010) and Ross (2021b: §3) for an illuminating discussion of the use of DNA evidence in trials; consult this page to learn more about post-conviction DNA exonerations <https://www.innocenceproject.org/about/>.

<sup>10</sup>The generally widespread convergence of intuitions against statistical evidence and in favour of individualised evidence may be telling but it's far from decisive. An experimental study conducted by Ebert et al. (2018) suggests that laypeople attribute a positive epistemic status to lottery-style propositions supported by sufficiently high statistical evidence. Moreover, the preference for individualised evidence can be explained by appeal to cognitive bias: Pennington and Hastie's (1991) story model of trial predicts that jurors' decision making is guided by attempts to impose a narrative structure to the evidence presented in court. Crucially, the informational richness of individualised evidence (e.g., eye-witness testimony) displays the exact narrative structure that could explain why it's usually preferred to individualised evidence (see Ross 2020a for helpful discussion). I offer these remarks to weaken the weight that is sometimes attributed to people's intuitions against statistical evidence.

<sup>11</sup>Another empirical observation about the law that warrants caution. The point on sufficiency doesn't always apply: in some key civil cases like toxic torts, statistical evidence is deemed sufficient. This is made especially vivid by the doctrine of market share liability (see *Sindell v. Abbott Laboratories* for a *locus classicus*). Suppose that a defective product causes harm to the plaintiff and it's hard to establish which of the many manufacturers who sell it did in fact sell the defective product to the plaintiff. In this case, each manufacturer selling the product has to reimburse the plaintiff according to the market share which, crucially, is determined by statistics. See Ross (2021b: 9; 2021a) and Enoch and Fisher (2015: 562) for an illuminating discussion of the genuine doctrinal complications with the use of statistical evidence in actual case law.

<sup>12</sup>Thomson (1986), Enoch et al. (2012), Buchak (2014) Smith (2016: Chapter 2, 2018), Blome-Tillman (2015, 2017), Pritchard (2018), Littlejohn (Littlejohn), Moss (forthcoming,2021), Backes (2020), Gardiner (2019,2020), Ross (2020a, 2020b, 2021a, 2021b), Di Bello (2019), Fratantonio (2021). It's worth noting that the puzzle has also received a decision-theoretic discussion in Regan et al. (2001).

<sup>13</sup>Hedden and Colyvan (2019) offer a contrary view.

conditions on individualised evidence are posited *also* as anti-luck conditions on knowledge. In this paper, I am not providing a novel solution to the puzzle. Rather, I shall argue that there can't be one in terms of a single anti-luck condition on knowledge. But to do so, I first have to clarify the link between knowledge, luck and individualised evidence.

At closer inspection, such link is not surprising: knowledge and individualised evidence share several common features. Knowledge is factive, so let's start with truth. Perhaps controversially, we can assume that individualised evidence also retains an essential connection to truth: it's propositional and consists of true propositions. This is because one key subproject in legal epistemology (though not the only one) is a *normative* project concerned with the type of evidence that an *ideal* legal system would employ.<sup>14</sup> Once we appreciate such normative dimension and concede that accounts of individualised evidence are cast as normative theses, it should be easy to appreciate also that an ideal system would not (and should not) base convictions on misleading evidence.<sup>15</sup>

Like knowledge, individualised evidence requires more than truth. The point is familiar: it is an important datum in epistemology that knowledge and luck are essentially incompatible.<sup>16</sup> The same holds in the law: legal verdicts which are correct only as a result of luck violate fundamental principles of justice. To appreciate why, recall the thought experiment offered by Judith Thomson (1986: 213). Imagine an incredibly lucky jury that ends up convicting the guilty by basing every verdict on a series of coin flips. Intuitively, the correctness of the verdicts does not guarantee fairness. Legal fact-finders are not concerned with mere accuracy: plausibly, they also aim to instil public confidence in the trial as an institution. They must ensure not only that the verdicts are right and just, but also that they are *perceived* to be right and just. To do so, an anti-luck condition is needed: just like knowledge is incompatible with luck, so is

---

<sup>14</sup>Some legal epistemologists explicitly engage in a normative project (e.g., Hedden and Colyvan 2019: 449, 458). Others are more interested in the descriptive enterprise of interpreting existing legal practice, and prefer to remain neutral on the potentially revisionary consequences of their views (e.g., Smith 2018: 1200).

<sup>15</sup>I must add an important point of clarification. This paragraph won't be palatable to those who reject a factive standard, and see accommodating the possibility of (procedurally correct) wrongful convictions as a virtue (Smith 2018: 1202; Gardiner *forthcoming*: 7). Perhaps more sophisticated versions of factive standards can be reconciled with the possibility of wrongful convictions (e.g., Blome-Tillman 2017), but a full elaboration of this claim exceeds the scope of the paper. To the unconvinced reader, I want to offer the following remarks. First: the necessary conditions discussed in this paper *inevitably* imply truth. Causal relations do not obtain without causal relata, sensitivity implies factivity (false beliefs are never sensitive), and so does the counterfactual formulation of safety originally offered by Sosa (1999: 146), according to which if S were to believe *p*, then *p* would be true ( $B(p) \rightarrow p$ ). Second: as the opening quote from Thomson shows, we have good independent reasons to think that truth, even if necessary, is ultimately insufficient for individualised evidence (hence the need for an anti-luck condition). In light of these points and in order for my discussion to get off the ground, I ask the reader to charitably concede that individualised evidence requires truth.

<sup>16</sup>See Unger (1968) for early work on luck, and Carter and Berrocal (2017) for a more recent overview.

the evidence upon which a legal verdict is based.

In light of these common features, striking similarities between the project of analysing knowledge and the attempt to provide an epistemic account of individualised evidence are brought into focus. According to Goldman (1967), knowledge requires causal relations, and so does individualised evidence according to Judith Thomson (1986). Nozick (1981) argued that knowledge requires sensitivity to truth, and so does individualised evidence according to Enoch, Fisher and Spectre (2012). Sosa (1999), Pritchard (2005, 2012) and Kelp (2013) have proposed a safety condition on knowledge: other legal epistemologists – Pritchard included – have followed suit and imposed a safety condition on individualised evidence (Pritchard 2018, Pardo 2018).

I now want to pause to register the following: the two debates significantly mirror each other. Within traditional epistemology, these conditions were proposed to fix the tripartite analysis of knowledge (allegedly) refuted by Gettier.<sup>17</sup> Importantly, an adequate solution to the Gettier problem presupposes a successful anti-luck condition: after all, the tripartite analysis clearly lacked one, as Dancy (1985: 134) and Pritchard (2015: 94) have later observed. Within legal epistemology, the same conditions were proposed to identify individualised evidence, which, like knowledge, also presupposes a similar non-accidental link to truth. This is the structural connection between knowledge and individualised evidence, and legal epistemologists have taken notice of it: more or less explicitly, they have assumed that the very same anti-luck condition on knowledge will also be necessary for (and so properly characterise) individualised evidence. Redmayne (2008: §4) dubs such an approach “knowledge-based” approach.

The knowledge-based approach enjoys initial plausibility, but I worry that it is not being pursued correctly (perhaps it’s not even worth pursuing; more on this in §5). Let’s pay heed to a key neglected detail. It bears noting that the knowledge-based approach proceeds in a somewhat *indirect* fashion: rather than on knowledge, so far the action has been on a *separate* anti-luck condition on knowledge. Building on this, I want to highlight an important and yet hitherto overlooked feature of the knowledge-based approach. The structural connection between knowledge and individualised evidence has been understood in a very *indirect* fashion: in fact, the focus is on separate anti-luck conditions on knowledge like (appropriate) causation, sensitivity and safety. The knowledge-based approach assigns explanatory priority to these separate causal or modal anti-luck conditions on knowledge and, as such, it turns out to be an

---

<sup>17</sup>See Ichikawa and Steup (2018), and Antognazza (2015) for a more accurate historical reconstruction.

essentially *indirect* approach.

The indirect way in which legal epistemologists have developed the knowledge-based approach deserves special scrutiny. The assumptions underlying it are suspicious, and rest on shaky foundations. These separate anti-luck conditions on knowledge featured in once prominent analyses, but, after all, the project of analysing knowledge has not borne much fruit. As Williamson (2000: 2-5, 30) observes, given the endless plethora of counterexamples faced by every purported analysis of knowledge, it is reasonable to conclude that none of them is problem-free. This invites the following question: granting that there is a structural connection between knowledge and individualised evidence, why think that such connection is best couched in terms of problematic anti-luck conditions on knowledge? Once imposed on individualised evidence, these conditions are likely to face the same difficulties that they face with knowledge. And this is precisely what we find. If legal epistemologists aim to uncover the structural connection between knowledge and individualised evidence, they'd better learn from the problems of anti-luck conditions on knowledge in lieu of simply reimposing them on individualised evidence.

In the remainder of the paper, I will substantiate this thought and show that, once imposed on individualised evidence, prominent anti-luck conditions on knowledge run into the same familiar and perhaps unavoidable difficulties that we see in post-Gettier epistemology. My strategy paves the way for a potentially far-reaching conclusion, which I revisit more extensively in §5: if I am right, no single causal or modal anti-luck conditions on knowledge can succeed as the right anti-luck condition on individualised evidence. As I will clarify, this conclusion has wider repercussions not only for the relevance of traditional epistemology to the understanding of individualised evidence, but also for legal epistemology more in general.

Having clarified these issues, I now proceed to assess three prominent and indirect knowledge-based approaches, starting with Thomson's causal account of individualised evidence.

### §3 Causation

Two elements stand out in ACCIDENT and PRISONERS. Firstly, statistical evidence lacks the right type of *direct connection* with the relevant facts. Secondly, statistical evidence does not eliminate the possibility of luck: impositions of liability and guilt based on statistical evidence can still be correct as a matter of luck. Judith Thomson (1986) takes these observations as starting points, and develops an account of individualised evidence in terms of *causal relations*.

Her account comprises two theses. First, individualised evidence plays an essential role:

it provides a non-accidental *guarantee* of guilt and liability. Second, only causally specified evidence provides such guarantee. Putting the two together, Thomson concludes that individualised evidence is causally specified: causal relations provide a guarantee which is meant to be immune to malign luck. Crucially, the resemblance with Goldman's (1967) causal analysis of knowledge is stark:

**Causal Analysis of Knowledge.** S knows that  $p$  if and only if S's belief that  $p$  stands in an appropriate causal relation to the fact that makes  $p$  true.

**Causal Individualised Evidence.** Individualised evidence guarantees that  $p$  if and only if it stands in an appropriate causal relation to the fact that makes  $p$  true.

In keeping with the indirect knowledge-based approach, Thomson moves from an anti-luck causal condition on knowledge to an anti-luck causal condition on individualised evidence. Her account is motivated by the overarching role that luck plays in making a verdict of guilt based on statistical evidence only accidentally correct (Thomson 1986: 214), so I'd better be clear on how luck relates to statistical and individualised evidence.<sup>18</sup>

Following Pritchard (2012), I read Thomson as holding that individualised evidence, unlike statistical evidence, should be incompatible with *veritic* luck, a type of luck which makes a belief only accidentally true and thus prevents it from constituting knowledge. Veritic luck operates at different levels. In the well-known barn façade example originally credited to Gettier, veritic *environmental* luck makes a belief luckily true for reasons having to do with the external environment: the belief is true but it could easily have been false. In standard Gettier cases (for instance, the sheep in the field case discussed in Chisholm 1966: 105), veritic *intervening* luck interferes in a more direct manner: an otherwise false belief accidentally hits the truth only courtesy of an intervention of luck. In what follows, I assume that both types of veritic luck are incompatible with individualised evidence (perhaps illegitimately: I return to this below).

We now get a better purchase on Thomson's causal account: while statistical evidence is compatible with veritic luck, causal individualised evidence is importantly immune to it. However, her proposal is plagued by the same problems that beset a causal theory of knowledge, and one of such problems has to do precisely with veritic luck. *Pace* Thomson, causal relations are compatible with environmental luck, as we learned from Goldman's early causal theory of knowledge. Consider the following case:

---

<sup>18</sup>We should grant to Thomson that causation is not probabilistic: if paired with more recent probabilistic theories of causation (see Hitchcock 2021: §2), her account of individualised evidence would look very different.

FAKE KNIVES. The famous detective Sherlock Holmes and his loyal sidekick John Watson are trying to solve a murder case. The manner of death of the victim is fairly standard (stabbed to death multiple times with a knife), but the crime scene is peculiar: it's filled with countless *seemingly identical* bloody knives. In order to throw off the detectives, the murderer made sure not only to cover each knife with the blood of the victim, but also with the fingerprints of random innocent people. Crucially, only one of them was used as the actual murder weapon – it's been used to kill the victim and it bears the fingerprints of the murderer. Sherlock inspects the scene, and by sheer luck he picks the very knife which was used as the murder weapon. Watson brings the knife as hard evidence to trial, and the murderer is convicted on the basis of the incriminating bloody knife.

This is a fake-barn case directed against Thomson. Its correct description is the following. The evidence is individualised and causally specified, and appropriate causal relations obtain between the relevant facts. The bloody knife is present on the crime scene *because* it was used as the murder weapon; the bloody knife bears the fingerprints of the murderer *because* they used it to commit the crime; the murderer is guilty *because* they used the knife to commit the crime and *because* the knife bears their fingerprints; and so on. However, just like the risk of error is too high to secure knowledge in fake-barn county, so the risk of error is too high in the case of a conviction based on this piece of causally specified individualised evidence. The verdict is only luckily correct. It could have easily been mistaken, and based on a seemingly identical knife which would have further led to the wrongful conviction of an innocent. Clearly, the verdict is correct only courtesy of a double stroke of luck: one bad (there are many identical murder weapons), one good (Sherlock happens to pick the right one). Causal individualised evidence is compatible with undermining environmental luck, and thus fails to provide a non-accidental guarantee.

I offered this case as a proof of concept that we can still have causally specified individualised evidence and malign veritic luck.<sup>19</sup> In order to better appreciate this point and to abstract away from too complicated thought experiments, I would like to offer two Gettier-style variations of stylised courts cases like ACCIDENT and PRISONERS:

FAKE BLUE BUSES. While driving back home, a vehicle suddenly hits Jane's car. In the accident, she breaks her leg. Thankfully, a passer-by witnesses the event: it was a blue bus that hit Jane's car, and so does the witness reports. However, there's a twist to the story: many more buses painted in blue but actually driven by the competing red bus company

---

<sup>19</sup>You may feel discomfort at advancing the debate on the basis of such seemingly contrived cases. I do too: I return on this in §5 and point out this fundamental methodological problem for the knowledge-based approach.

drivers were also present on the scene of the accident and also operated a route at the same time of the event. The eyewitness correctly identifies the only genuine and liable blue bus, but they do so only as a result of sheer luck.

TWIN PRISONERS. 100 prisoners are exercising in the prison yard. Suddenly 99 of them attack the guard, putting into action a plan that the 100th prisoner knew nothing about. The 100th prisoner played no role in the assault and could have done nothing to stop it. Thankfully, a guard witnesses the event, and correctly identifies prisoner number 58 as guilty. However, there's a twist to the story: prisoner 58 has an identical twin brother who happens to be exactly the 100th innocent prisoner. The eyewitness correctly identifies the guilty prisoner, but they do so as a result of sheer luck.

To repeat: these are further fake-barn variations of ACCIDENT and PRISONERS. Here too, the evidence is individualised and causally specified, yet it fails to exclude veritic luck due to misleading features of the scene. Again, appropriate causal relations obtain between the relevant facts. The blue bus company is liable because a blue bus hit Jane's car, and the passer-by saw a blue bus hitting Jane's car because a blue bus did in fact hit Jane's car. Prisoner 58 is guilty because they assaulted the guard, and the eyewitness saw prisoner 58 assaulting the guard because prisoner 58 did in fact assault the guard. However, even if based on causally specified individualised evidence, the verdict is only luckily correct: the passer-by could have too easily looked at a fake blue bus driven by a red bus driver, and then the piece of individualised evidence would have incorrectly determined that the blue bus company is liable. Similarly, the eye-witness could have too easily mistaken the guilty prisoner 58 for their innocent identical brother, and then the evidence would have incorrectly determined that prisoner 100 is guilty. The point remains: causally specified individualised evidence is still compatible with an only (veritically) luckily correct verdict.

We have now come full circle. Thomson argues against statistical evidence because verdicts based on it can be correct as a matter of luck, and she goes on to impose a causal anti-luck condition on individualised evidence by leaning on a causal theory of knowledge. However, her proposal suffers from the same problem she raises, since causal relations are equally compatible with veritic environmental luck.<sup>20</sup> Tellingly for my purposes, my objection against her causal account mirrors the counterexample against the causal analysis of knowledge.

---

<sup>20</sup>There is increasing tendency to accept that propositional knowledge is compatible with environmental luck, most notably in Sosa's bi-level virtue epistemology (Sosa 2015: 77-81). I remain neutral on whether propositional knowledge is compatible with environmental luck: it should be granted that agents run a high possibility of error in these scenarios, and that causally specified evidence does not insure against environmental luck. This suffices to raise the objection against causally specified evidence, independently of the potentially more controversial

## §4 Modal Conditions

The same problem affecting a causal theory of knowledge also spells trouble for a causal account of individualised evidence. Causal relations are compatible with veritic luck, and therefore fail as anti-luck condition on both knowledge and individualised evidence. One way of diagnosing this problem is to hold that individualised evidence should appropriately *track* the truth; another is to maintain that it should not *too easily* support a falsehood. These observations pave the way for two modal anti-luck conditions on knowledge, sensitivity and safety. Can they do better than causal relations as anti-luck conditions on individualised evidence? In the following two subsections, I answer negatively. In §4.1, I show that sensitive evidence also falls prey to veritic luck, and point out further difficulties with other canonical instances of individualised evidence. Next, in §4.2., I argue as follows. Even granting that safety is a successful anti-luck condition, the principle does not carve at the right epistemological joints: in fact, it fails to distinguish between statistical and individualised evidence.

### 4.1 Sensitivity

The sensitivity condition on knowledge was introduced by Robert Nozick, who suggests that if S knows that  $p$ , then S's belief that  $p$  tracks the truth. Following Nozick (1981), I break down the sensitivity principle into two twin conditions, variation and adherence. Here is a full-blown statement of sensitivity:

**Sensitivity.** S knows that  $p$  if and only if: had  $p$  been false, then S would not have believed  $p$  (variation); and, had  $p$  been true, then S would have believed  $p$  (adherence).

Enoch, Fisher and Spectre (2012) build on Nozick's analysis of knowledge and impose a sensitivity condition on individualised evidence.<sup>21</sup> Accordingly:

**Sensitive Evidence.** Individualised evidence  $e$  is sensitive to  $p$  if and only if: had  $p$  been false, then S would not have believed  $p$  on the basis of  $e$  (variation); and, had  $p$  been true, then S would have believed  $p$  on the basis of  $e$  (adherence).

---

claim that knowledge and environmental luck are incompatible (e.g., Pritchard 2012: §3). Notice further that causally specified evidence does not insure against intervening luck either: I could equally press a case of deviant causal chains against Thomson, but here I've decided to focus on environmental luck to highlight the analogy with Goldman's causal theory of knowledge.

<sup>21</sup>To forestall misunderstanding, I flag that sensitivity plays only a *partial* role in Enoch and colleagues' overall account of individualised evidence. While they assume a sensitivity account of individualised evidence, to avoid what they call knowledge-fetishism their incentive-based explanation carries more theoretical weight than sensitivity (Enoch et al. 2012: 220-223). But sensitivity remains relevant to their incentive-based explanation: they acknowledge that "the epistemic story and the incentive-based story are closely knit and interestingly related in light of their similar structures and ramifications in the legal arena" (Enoch and Fisher 2015: 565).

With an eye on the problems encountered by a sensitivity condition on knowledge, I shall develop two main foci of attack against sensitive evidence, which is either too weak or too strong. Too weak: it fails to safeguard against veritic luck. Too strong: it rules out insensitive but acceptable instances of individualised evidence. Let me take each point in turn.

First off, sensitivity does not preclude (Gettier-style) intervening luck. In fairness to sensitivity theorists, their principle wasn't originally offered as an explicit anti-luck condition. However, several authors - Nozick included - have conceded that sensitivity should preclude standard Gettier cases, which in turn involve intervening veritic luck.<sup>22</sup> Sensitivity does not meet this desideratum: more ingenious Gettier cases can be raised with equal force also against sensitivity.<sup>23</sup> Here is a particularly apt one:

LUCKY COLD HIT. A team of detectives is investigating a murder. The evidence is scarce, but some biological traces are available. Genetic material is found at the crime-scene: it's a strand of hair belonging to the perpetrator. It gets tested, and the results bring in a new suspect, Mr. X, who has indeed committed the crime. Upon further investigation, Mr. X is found guilty and convicted on the basis of the DNA evidence. There is, however, a twist to the story. The DNA found on the scene matched the new suspect only by luck: courtesy of a previous lab error, Mr. X's long deceased monozygotic twin brother, Mr. Y, was accidentally registered in the DNA database under the name of Mr. X. While it's easy to see why the mistake was made (their DNA profile is identical), the lab error is nevertheless the chief reason why the DNA database included Mr. X's profile in the first place.

The case I offer is a Gettierised cold-hit DNA match. As Smith (2018: 1215) also notes, DNA evidence is usually sensitive, and satisfies sensitivity's two-fold condition. In LUCKY COLD HIT, the evidence satisfies variation. Had Mr. X not committed the murder, he wouldn't have left any genetic material at the crime scene and there would have been no hit. On such basis, the investigators *would not* have believed him to be guilty. Likewise, the evidence satisfies also adherence. Had Mr. X committed the murder in relevantly similar circumstances, he would have left genetic material at the crime scene, and the hit would have taken place. On such basis, the investigators *would have* still believed him to be guilty. Yet in this case sensitive DNA evidence tracks the truth and supports a correct verdict only courtesy of intervening luck: the DNA belongs to the culprit, but the hit is lucky and the match accidental. In this respect, sensitivity is too weak to provide a successful anti-luck condition.

---

<sup>22</sup>Pritchard (2005: 158; 2013: §2) discusses the role of sensitivity in handling Gettier cases. Nozick (1981: 173) admits that his sensitivity condition should serve as anti-Gettier condition.

<sup>23</sup>See also Goldberg's MOVIE SET (Goldberg 2015: 279), and Becker's VASE HOLOGRAM (Becker 2018: 124).

In another respect however, sensitivity is too stringent: if knowledge required sensitivity, we'd have too little of it. Similarly, once sensitivity is imposed as necessary condition, the resulting conception of individualised evidence is extremely impoverished. Sensitivity theorists struggle to account for cases of safe but insensitive inductive knowledge, as Sosa (1999: 145) first noticed. In like manner, pieces of inductive evidence are safe but equally insensitive. Consider:

KANTIAN EYE-WITNESS. Maria likes to take strolls. She does this with extreme reliability, and her habit is often compared to Kant's famous afternoon walks. During her strolls, Maria regularly walks on the street where the accident takes place, and regularly sees a blue bus operating the route. The day comes, and a blue bus hits the car. Maria doesn't witness the event, but testifies that she spotted a blue bus shortly before the time of the crash, and that she has seen a blue bus operating the route in that very street every other day until that day. To back up her statement, she lists every time she's taken her walk and seen a blue bus operating the route around the time and place of the crash. The list is long, and provides (inductive) reason to conclude that a blue bus is liable.

This is a case of inductive eye-witness testimony (and, *a fortiori*, individualised evidence). Sensitivity gets the wrong result here: Maria's testimony may be insensitive, but it's nevertheless acceptable in virtue of being safe. Her testimony is insensitive: in the closest possible world where a blue bus is not liable (perhaps because a red bus mysteriously went off-track and hit the car), one forms a false belief on the basis of her testimony. Still, her testimony is safe, and thus acceptable: in more relevant close worlds, a blue bus is liable, and one *does* form a true belief on the basis of her testimony. Because of sensitivity, this key aspect simply goes overlooked.

To further appreciate, on more general grounds, why sensitivity is a too strong necessary condition for individualised evidence there's one final case I want to offer - an epistemic Frankfurt case. Notably, such cases have been successfully deployed against safety conditions (Kelp 2016), but they apply with equal (if not major) force also against sensitivity. Consider:

SCHEMING PRISONER. 100 prisoners are exercising in the prison yard. Suddenly 99 of them attack the guard, putting into action a plan that the 100th prisoner knew nothing about. Prisoner 54 is guilty, but another scheming prisoner is not aware of this and wants to make sure that prisoner 54 is punished anyway due to long-standing resentment against him. The scheming prisoner hatches his plan to frame the assault on prisoner 54, fully ready to tamper with any piece of exculpatory evidence. However, he finds an authentic videotape that already implicates prisoner 54 and, happy with the result, he doesn't intervene.

Intuitively, the videotape provides insensitive but acceptable individualised evidence of guilt.

The evidence is insensitive: in the closest possible world where prisoner 54 is innocent, the scheming prisoner intervenes, tampers with the evidence (the videotape) and frames him. Had  $p$  been false, one would have believed  $p$  anyway given the evidence. More precisely: had prisoner 54 been innocent, one would have believed him to be guilty anyway thanks to the suitable intervention of the scheming prisoner. Yet, even if insensitive, the evidence is acceptable. Any reason to discard it seems unmotivated: unlike standard Gettier and fake-barn cases, there's no *actual* environmental or intervening luck involved. There's no deviant causal chain, and no clear indication of any salient impropriety (epistemic or otherwise). Because of sensitivity, we are forced to deny the intuitive status of acceptable individualised evidence to the videotape. The lesson to learn here is general: sensitivity is too strong a condition for individualised evidence.<sup>24</sup>

Let me take stock. I've developed a two-pronged attack against sensitive evidence. Firstly, sensitivity is too weak to safeguard against intervening luck: Gettier-style cases are easy to find and generate (LUCKY COLD HIT). Secondly, sensitivity is also too strong: it rules out safe but insensitive inductive evidence (KANTIAN EYE-WITNESS), and delivers the wrong verdict on plain cases of acceptable but insensitive individualised evidence (SCHEMING PRISONER). Here too, problems faced by a sensitivity condition on knowledge apply also to sensitive evidence.

## 4.2 Safety

Neither causation nor sensitivity succeed in excluding veritic luck. One final candidate left to consider is another anti-luck condition proffered by epistemologists, the safety principle. Pritchard also moves from a safety condition on knowledge to one on individualised evidence:

The modal condition that we are imposing on legal evidence in this regard is what is known as a safety condition. Such a condition is generally regarded as being a necessary condition for knowledge, and also – relatedly – being the condition that excludes the kind of epistemic luck/risk which is incompatible with knowledge. (Pritchard 2018: 117)

First, it bears noting that several authors have questioned the success of safety as anti-luck condition: a true belief can be lucky even if formed through a safe method.<sup>25</sup> Thus, *pace*

---

<sup>24</sup>One may worry that epistemic Frankfurt cases put pressure on *any* modal anti-luck condition on knowledge and individualised evidence (both sensitivity and safety). After all, such cases were originally designed to show that safety - rather than sensitivity - is a too strong condition (Kelp 2016: §2.1). The reader is free to draw this stronger conclusion, but I aim for less. I offer the Frankfurt case to show that sensitivity is too strong, and I prefer to remain non-committal on different (and perhaps less clear-cut) verdicts pertaining to safety. As I will show in the next subsection, safety accounts of individualised evidence are problematic on independent grounds.

<sup>25</sup>See for instance Hiller and Neta (2007: 208); Lackey (2006: 288); Coffman (2010: 246); Goldberg (2015: 275). To appreciate the strength of such counterexamples, it's worth emphasising that while safety is offered as a merely necessary condition on knowledge, it is nevertheless meant to be sufficient to exclude knowledge-undermining luck

Pritchard, safety may not succeed to exclude the kind of epistemic luck which is incompatible with knowledge. However, the problem is worse than this: even if safety were a successful anti-luck condition, it would not draw the line in the right place. In fact, in both its weak and strong formulations, safety does not distinguish between statistical and individualised evidence.

Pritchard argues that while statistical evidence is unsafe, individualised evidence is safe (Pritchard 2018: §3). This is because, on the basis of the former, the modal risk of forming a false belief is high: little would have to change in the actual world for a belief based on statistical evidence to be false. Statistical evidence is *unsafe*: in similar nearby worlds, beliefs based on it are false. Conversely, on the basis of the latter, the modal risk of forming a false belief is low: a lot would have to change in the actual world for a belief based on individualised evidence (e.g., eye-witness testimony) to be false. Individualised evidence is *safe*: in similar nearby worlds, beliefs based on it continue to be true. These are the core claims of Pritchard’s safety account.

Pritchard does not commit himself to a specific formulation of safety. Thus, the success of his safety account turns on a preliminary terminological disambiguation and two distinct questions. First, we need to get clear on whether safety is understood weakly or strongly. Having clarified this, we can properly assess the safety account by focussing on the following two questions. Given a more refined formulation of the safety principle, is Pritchard right in claiming that statistical evidence is unsafe (Question 1) and that individualised evidence is safe (Question 2)? Let’s begin with the weaker formulation of safety:

**Weak Safety.** S believes  $p$  on the basis of safe evidence if and only if S’s belief that  $p$  is true in the actual world and in *most* close possible worlds in which S forms the belief that  $p$  on the basis of the same evidence S uses in the actual world,  $p$  is true.<sup>26</sup> (Greco 2012: 196)

Statistical evidence satisfies weak safety. Recall our initial examples: let’s grant that prisoner 54 is guilty and that a blue bus hit Jane’s car. These are not random aleatory events: as Gardiner (2020: §6) notices, guilty prisoners may riot also in close worlds. Ditto for liable bus drivers, who may equally drive negligently in close worlds too. Statistical evidence can be weakly safe: there can be most close worlds where prisoner 54 continues to assault the guard, and a blue bus still hits Jane’s car. Thus, on its basis, one does form a true belief in most close worlds.

---

(Pritchard 2013: 158).

<sup>26</sup>While standard construals of safety assume truth, there’s an additional reason to stipulate that  $p$  is true: the truth of  $p$  impacts on its modal profile. As Ebert et al. observe (2018: 10-11), no world is more similar to the actual world than the actual world itself. It’s thus difficult to make assessments on safety without already taking a stand on whether the belief in question is true or false or on whether the defendant is guilty or innocent. This is a structural feature of Pritchard’s conception of safety and risk, as Smith (2018: 1204-1205) also notices.

The vignettes are admittedly underdescribed, but nothing in their description suggests that the modal profiles of events like a riot or an accident are fragile: plausibly, prisoner 54 does not just riot on a whim, nor does the bus driver “just so happen” to hit Jane’s car. Once these events obtain in the actual world, they also plausibly obtain in most close worlds.<sup>27</sup>

The foregoing points provide an answer to Question 1. In the cases of interest, beliefs based on sufficiently strong statistical evidence are weakly safe: contrary to Pritchard’s account, in nearby worlds beliefs based on it continue to be true. The prospects for the safety account may start to look dim by now, but not all is lost for the safety theorists, who can still appeal to a stronger formulation of safety:

**Strong Safety.** S believes  $p$  on the basis of safe evidence if and only if S’s belief that  $p$  is true in the actual world and in *every* close possible world in which S forms the belief that  $p$  on the basis of the same evidence S uses in the actual world,  $p$  is true. (Greco 2012: 196)

Strong safety is a more promising principle for ruling out statistical evidence, which fails to satisfy this stronger requirement. After all, there will be one or more close possible worlds where, on the basis of statistical evidence, one will falsely believe that prisoner 54 attacked the guard, or that a blue bus hit Jane’s car. Accordingly, strong safety gives a compelling negative answer to Question 1. So far, so good. What about Question 2? Here, safety theorists have additional work to do: they have to show not only that statistical evidence is not (strongly) safe, but also that individualised evidence is (strongly) safe. However, under a strong formulation of the safety principle this is not feasible: just like statistical evidence, individualised evidence also fails to satisfy a stronger safety requirement.

Here is why. Take a canonical piece of individualised evidence such as eye-witness testimony. Inevitably, there will be a close possible world where one forms a false belief on its basis. All we need is a close enough world where a very improbable possibility is actual. Consider:

**DISHONEST EYE-WITNESS.** While driving back home, a vehicle suddenly hits Jane’s car. She realises that the vehicle was a bus, but she cannot discern the colour. A huge crowd of bystanders witness the event. Jane asks to one of them the colour of the bus which hit the car. The bystander answers sincerely: it was a blue bus. Every other bystander would have answered in the same way – all except one. In fact, unbeknownst to Jane, among the

---

<sup>27</sup>Safety theorists can respond to Gardiner’s challenge in (at least) three ways. First, they may strengthen the safety principle and deny that statistical evidence is safe. However, as I explain below, this is a costly answer. Second, they may index the safety principle to specific belief-formation methods. This is also a costly answer, likely to raise the generality problem famously pressed by Conee and Feldman (1998). Third, they may restrict the focus to normal worlds instead of similar close worlds (see Smith 2016: Chapter 6 for a comparative discussion of safety and normalcy). I take this to be the best response, but I lack the space to consider it fully here.

bystanders there's also a dishonest eye-witness, who works for the blue bus company and is prepared to lie to cover for his company. Luckily, Jane picks one of the many sincere bystanders instead of the only dishonest eye-witness.<sup>28</sup>

This is our close but improbable possibility. Jane could easily have formed a false belief on the same basis: there's a very similar close possible world where she casually picks the dishonest eye-witness instead of the sincere bystander. Since strong safety does not tolerate the possibility of error in *any* close world, this canonical instance of individualised evidence is not strongly safe. Again, contrary to Pritchard's account, in nearby worlds beliefs based on it are false.

Overall, strong safety places too demanding a requirement on individualised evidence, which, plausibly enough, supports the truth of a proposition in *most* close worlds rather than in *every* close world. As a necessary condition, strong safety comes at a high cost: while it may rule out statistical evidence, it excludes also canonical pieces of individualised evidence.

At this point, each of Pritchard's core claims is cast into doubt. Since it satisfies weak safety, it's not the case that statistical evidence is unsafe. And, since it does not satisfy strong safety, it's not the case that individualised evidence is safe. Safety theorists now face a dilemma: either individualised evidence is weakly safe [*Option 1*], or it is strongly safe [*Option 2*].<sup>29</sup> Suppose that they choose [*Option 1*]: if individualised evidence is weakly safe, then so is statistical evidence. Weak safety fails to distinguish between the two. Suppose that they choose [*Option 2*] instead: if individualised evidence is *strongly* safe, then neither statistical evidence nor canonical instances of individualised evidence (e.g., eye-witness testimony) are strongly safe. Strong safety also fails to distinguish between the two.

This is a minimum desideratum for a successful safety-based account: individualised evidence should have a modal epistemic quality (weak or strong safety) that statistical evidence lacks.<sup>30</sup> Safety does not deliver on this front. No matter which option is picked, safety does not carve at the relevant epistemological joints: it does not distinguish between statistical and individualised evidence. It looks like the right anti-luck condition on individualised evidence just can't be found; I now proceed to say more on why it can't be found.

---

<sup>28</sup>To keep the discussion manageable, I ask the reader to bracket complicated issues pertaining to further defeating and corroborating evidence Jane might acquire during later stages of her inquiry.

<sup>29</sup>Safety-theoretic accounts of knowledge face similar objections (e.g., Greco 2007: §1; Sosa 2015: 119; Gardiner 2020: §6). The dilemma I raise here is a hybrid and hence novel one: in [*Option 1*], I borrow elements from Gardiner's objection. In [*Option 2*], I draw on Greco's and Sosa's versions.

<sup>30</sup>Blome-Tillman (2015) also implicitly accepts this desideratum. He objects to the sensitivity and causal accounts of individualised evidence by showing that statistical evidence can also be sensitive or causally specified.

## §5 Diagnosis

At this point, let's focus on the bigger picture: no account of individualised evidence in terms of single causal or modal anti-luck conditions on knowledge is or can be correct. Given the failure of these conditions on knowledge, we should expect them to fail also when imposed on individualised evidence. The expectation is now confirmed: these conditions fail on *both* knowledge and individualised evidence. But what should we make of this double failure? I suggest to make an inference to the best explanation: just like any analysis of knowledge in terms of separate causal or modal anti-luck conditions is unsuccessful, so will be any account of individualised evidence in terms of the same anti-luck conditions. Here's my diagnosis: no single causal or modal anti-luck condition on knowledge can succeed as anti-luck condition on individualised evidence. This also calls into question the relevance of these anti-luck conditions to the understanding of individualised evidence and to legal epistemology in general.

Before looking at the broader consequences of this diagnosis for legal epistemology, it's worth dwelling on its independent plausibility. That veritic luck is hard to exclude is a familiar take-home lesson from the post-Gettier literature: purported causal and modal conditions on knowledge fail to capture the right non-accidental connection to truth, nor insure against veritic luck. The problem is general: as Zagzebski (1994: 72) explains, veritic luck strikes "as long as there is a small degree of independence between truth and the other conditions on knowledge".

Given the inescapability of veritic luck, single causal or modal anti-luck conditions on knowledge are intrinsically defective. This problem is also general. Here's a telling quote from Goldberg:

No purely modal condition (of the sort epistemologists have used to analyze knowledge) will suffice as the anti-luck condition on knowledge, since we can concoct examples in which it is simply a matter of luck that the method has that modal property. (Goldberg 2015: 279)

So much for knowledge. However, and crucially for my purposes, this problem runs deeper. In fact, as explained in §3, the knowledge-based approach championed by some legal epistemologists takes knowledge and individualised evidence to be structurally connected by a similar anti-luck condition: Goldberg's pessimistic remarks apply also to individualised evidence. For sake of vividness, let me make this last point from a different angle. Consider this rough and ready equation, where  $x$  denotes a causal or modal anti-luck condition distinct from knowledge:

$$\text{Knowledge} \equiv \text{Truth} + \text{Belief} + x$$

These days, this equation has fallen out of favour. We learn from post-Gettier epistemology that while truth and some separate causal or modal anti-luck condition may be individually necessary, they're nevertheless jointly insufficient to exclude veritic luck and obtain knowledge. Despite many attempts, no uncontroversial  $x$  has been found. Suppose that, drawing on a strong inductive basis, we throw in the towel and follow Williamson (2000: 3-4) in thinking that no  $x$  can be found.<sup>31</sup> Then the same inductive basis gives us reason to reject also this rough and ready equation, where  $x$  still denotes an anti-luck condition distinct from knowledge:

$$\text{Individualised Evidence} \equiv \text{Truth} + x$$

Proponents of the knowledge-based approach accept or assume some version of this equation.<sup>32</sup> In this paper, I have considered three prominent candidates for  $x$ : a causal and two modal anti-luck conditions. However, just like these conditions fail to identify the non-accidental connection to truth which is necessary for knowledge, likewise they fail to identify the non-accidental connection to truth which is taken to be necessary for individualised evidence. What is more, they fail for similar reasons: objections to separate causal and modal anti-luck conditions on knowledge apply also to accounts of individualised evidence in terms of the same anti-luck conditions. Crucially, these objections are inescapable: they generalise to any account of individualised evidence built around a causal or modal anti-luck condition distinct from knowledge itself.

These points are significant for legal epistemology. The failure of single causal or modal anti-luck conditions on both knowledge and individualised evidence suggests that the notion of individualised evidence can't be successfully elucidated in terms of these single causal or modal anti-luck conditions. Because it employs these defective anti-luck conditions, the knowledge-based approach is unfruitful and fundamentally limited. Firstly, such approach consists in imposing on individualised evidence the very same anti-luck conditions that already fail when imposed on knowledge: as such, this way of proceeding appears to be a non-starter. Secondly, because it uses rather contrived counterexamples made popular in mainstream epistemology the knowledge-based approach fails to engage with the subtle and complicated realities of existing legal systems. The approach assumes a somewhat omniscient observer (standardly absent in

---

<sup>31</sup>See Cassam (2009: 21-22) and Williamson (2009: 290) for discussion on the unanalysability of knowledge.

<sup>32</sup>The equation is schematic, and does not represent (or attribute) a reductive analysis. I am concerned with anti-luck conditions on individualised evidence: while these conditions are meant to be sufficient to exclude luck, they are not meant to be sufficient for a positive verdict. Similarly, as noted above, anti-luck conditions on knowledge are meant to be sufficient to exclude knowledge-undermining luck, not sufficient for knowledge. I remain neutral on whether individualised evidence demands more than a successful anti-luck condition to suffice for a positive verdict. Even a fully successful anti-luck condition may be a necessary but still insufficient condition for this more demanding role.

legal courts), and, in the case of modal conditions, a somewhat precise and privileged access to the closeness of possible worlds that determine judgements on safety and sensitivity (a privileged access which is also standardly absent in legal courts). Thirdly, no matter which specific anti-luck condition is picked, anti-luck conditions in general fail to provide a meaningful distinction between statistical and individualised evidence. As we have seen, they do not draw the line in the right place: that there is no real distinction between statistical and individualised evidence remains a live option which the knowledge-based approach can't rule out. Accordingly, despite being motivated by what appears to be a common anti-luck condition on knowledge and individualised evidence, the knowledge-based approach is not very fruitful for legal epistemology. Given the focus on too many artificial cases, it is ill-suited to advance the understanding of legal evidence. And since it fails to provide a successful distinction between statistical and individualised evidence, it is also ultimately inconclusive. Taken together, these considerations suggest that as far as veritic luck is concerned, epistemology and the law seem to come apart.

This last point is best appreciated in the context of other important objections against the relevance of epistemology to the understanding of evidence law. The problems that single anti-luck conditions on knowledge incur when extended to individualised evidence speak in favour of Allen's complaint against a naive and simple application of epistemic concepts to the law. As he puts it (Allen 2021a: 255), this application often results in relying on "weird hypotheticals" that place "impossible epistemological demands" on a crucial type of legal evidence like individualised evidence.<sup>33</sup> The problems of the knowledge-based approach also bolster Fratantonio's recent case against a purely epistemic solution to the puzzle of statistical evidence (Fratantonio 2021). Borrowing from her useful taxonomy (Fratantonio 2021: §2), if legal epistemology is conceived as a descriptive project that aims at elucidating legal concepts (e.g., individualised evidence) in epistemic terms (e.g., causal or modal anti-luck conditions), then the main anti-luck conditions on knowledge fail to achieve this clarificatory aim: after all, they don't successfully distinguish between statistical and individualised evidence. And if legal epistemology is conceived as a normative project concerned with why an ideal legal system *should not* use statistical evidence, then the main anti-luck conditions on knowledge also fail to achieve this aim for similar reasons (we are still left without a meaningful distinction between statistical and individualised evidence). Overall, anti-luck conditions seem irrelevant for the pursuit of both a

---

<sup>33</sup>See Allen (2021a, 2021b, 2021c), Allen and Smiciklas (forthcoming) and Allen and Leiter (2001) for an illuminating elaboration of this criticism.

descriptive and a normative project in legal epistemology.

These remarks align well with further compelling attempts that build on normative considerations to cleave epistemology from evidence law. Since the knowledge-based approach has proven unsuccessful in tracking any substantial normative difference between statistical and individualised evidence, Enoch, Fisher and Spectre may be right in holding that epistemology is not “intrinsically and directly relevant to normative evidence law theory” (Enoch, Fisher and Spectre 2021: 85). Similarly, the limits of the knowledge-based approach provide further reason to think, as Ross (2021a) is keen to emphasise, that courts of law are subject to normative requirements fundamentally distinct from those applying to individual subjects routinely discussed in epistemology. There seems to be room for fruitful interactions between epistemology and theory of evidence law, but neither the knowledge-based approach nor the focus on anti-luck conditions are instances of such interactions.

I would like to note a few important differences with these influential objections to the relevance of epistemology to the understanding of legal evidence. Because it is circumscribed to anti-luck conditions, the point made here is more modest: unlike Fratantonio’s objection, the arguments in this paper do not target epistemic solutions to the puzzle of statistical evidence that do away with anti-luck conditions on knowledge. Similarly, unlike Enoch, Fisher and Spectre’s criticism, the arguments offered here don’t generalise to every epistemic notion. However, these arguments do retain the spirit of these influential points against the relevance of epistemology for legal evidence and also contribute to weaken such relevance.<sup>34</sup>

Having clarified this, I can take stock and conclude. In light of an intuitively similar anti-luck condition, one approach in legal epistemology has consisted in moving from an anti-luck condition on knowledge to one on individualised evidence. This *modus operandi* has been found wanting: no causal or modal anti-luck condition on knowledge can succeed as anti-luck condition on individualised evidence.

The chief moral to draw from post-Gettier epistemology is that causal and modal conditions on knowledge fail as satisfactory anti-luck conditions. If this much can be agreed upon, why

---

<sup>34</sup>Another conclusion consistent with this diagnosis is that the right anti-luck condition on individualised evidence is knowledge *simpliciter* rather than a separate modal or causal condition on knowledge. Since knowledge and veritic luck are by hypothesis incompatible, knowledge would (trivially) succeed as the right anti-luck condition on individualised evidence. I would however urge caution before drawing this conclusion. A flat-footed equation of knowledge and individualised evidence would place an even more demanding epistemological requirement on legal evidence, and it would do so on the basis an objectionable motivation (i.e., contrived versions of stylised court cases). Secondly, the putative difference between individualised and statistical evidence doesn’t seem to be knowledge *per se* but rather what explains absence of knowledge (e.g., appropriate causation, sensitivity or safety). A direct appeal to knowledge hardly makes progress on the problem.

should we expect them to succeed when imposed on individualised evidence? In this paper, my aim was to show that it is probably just a false expectation.<sup>35</sup>

## References

- Ronald Allen. Naturalized epistemology and the law of evidence revisited. *Quaestio facti. Revista internacional sobre razonamiento probatorio*, 0(2), 2021a. ISSN 2604-6202. doi: 10.33115/udg\_bib/qf.i2.22446. URL <https://revistes.udg.edu/quaestio-facti/article/view/22446>.
- Ronald Allen. Naturalized epistemology and the law of evidence revisited: Reply to enoch. *Quaestio facti. Revista internacional sobre razonamiento probatorio*, 0(2), 2021b. ISSN 2604-6202. doi: 10.33115/udg\_bib/qf.i2.22446. URL <https://revistes.udg.edu/quaestio-facti/article/view/22446>.
- Ronald Allen. Naturalized epistemology and the law of evidence revisited: Reply to pardo. *Quaestio facti. Revista internacional sobre razonamiento probatorio*, 0(2), 2021c. ISSN 2604-6202. doi: 10.33115/udg\_bib/qf.i2.22446. URL <https://revistes.udg.edu/quaestio-facti/article/view/22446>.
- Ronald Allen and Cristopher Smiciklas. The law's aversion to naked statistics and other mistakes. *Legal Theory*, forthcoming.
- Ronald J. Allen. Explanationism all the way down. *Episteme*, 5(3):320–328, 2008.
- Ronald J. Allen and Brian Leiter. Naturalized epistemology and the law of evidence. *Virginia Law Review*, 87(8):1491–1550, 2001. ISSN 00426601. URL <http://www.jstor.org/stable/1073901>.
- Ronald J. Allen and Michael S. Pardo. Relative plausibility and its critics. *The international journal of evidence & proof*, 23(1-2):5–59, 2019.
- Ronald J. Allen and Alex Stein. Evidence, probability, and the burden of proof. *Arizona law review*, 55(3):557, 2013.

---

<sup>35</sup>Many thanks to two anonymous reviewers for detailed comments on a previous version of the manuscript. Thanks also to Chris Kelp Adam Carter and Sven Bernecker for guidance and advice. I am especially grateful to Giada Fratantonio, Lewis Ross and Lilith Newton for many conversations about the topic of this paper. This work received generous funding from a Royal Institute of Philosophy Jacobsen scholarship.

- Maria Rosa Antognazza. The benefit to philosophy of the study of its history. *British Journal for the History of Philosophy*, 23(1):161–184, 2015. doi: 10.1080/09608788.2014.974020.
- Marvin Backes. Epistemology and the law: Why there is no epistemic mileage in legal cases. *Philosophical Studies*, 177(9):2759–2778, 2020. doi: 10.1007/s11098-019-01337-6.
- Gary S. Becker. Crime and punishment: An economic approach. *Journal of Political Economy*, 76(2):169–217, 1968. ISSN 00223808, 1537534X. URL <http://www.jstor.org/stable/1830482>.
- Kelly Becker. *The Sensitivity Response to the Gettier Problem*, page 108–124. *Classic Philosophical Arguments*. Cambridge University Press, 2018. doi: 10.1017/9781316827413.009.
- Marcello Di Bello. Trial by statistics: Is a high probability of guilt enough to convict? *Mind*, 128(512):1045–1084, 2019. doi: 10.1093/mind/fzy026.
- Michael Blome-Tillmann. Sensitivity, causality, and statistical evidence in courts of law. *Thought: A Journal of Philosophy*, 4(2):102–112, 2015. doi: 10.1002/tht3.163.
- Michael Blome-Tillmann. 'more likely than not' - knowledge first and the role of statistical evidence in courts of law. In Adam Carter, Emma Gordon, and Benjamin Jarvis, editors, *Knowledge First - Approaches in Epistemology and Mind*, pages 278–292. Oxford, UK: Oxford University Press, 2017.
- Fernando Broncano-Berrocal and J. Adam Carter. Epistemic luck. In Tim Crane, editor, *Routledge Encyclopedia of Philosophy*. Routledge, 2017.
- Lara Buchak. Belief, credence, and norms. *Philosophical Studies*, 169(2):1–27, 2014. doi: 10.1007/s11098-013-0182-y.
- Guido Calabresi. Some thoughts on risk distribution and the law of torts. *The Yale Law Journal*, 70(4):499–553, 1961. ISSN 00440094. URL <http://www.jstor.org/stable/794261>.
- Quassim Cassam. Can the concept of knowledge be analysed? In Duncan Pritchard and Patrick Greenough, editors, *Williamson on Knowledge*. Oxford University Press, 2009.
- Roderick Milton Chisholm. *Theory of Knowledge*. Englewood Cliffs, NJ, USA: Englewood Cliffs, N.J., Prentice-Hall, 1966.

- C. A. J. Coady. *Testimony: A Philosophical Study*. Oxford, England: Oxford University Press, 1992.
- E. J. Coffman. Misleading dispositions and the value of knowledge. *Journal of Philosophical Research*, 35:241–258, 2010. doi: 10.5840/jpr\emph{2010}13.
- Laurence Jonathan Cohen. *The Probable and the Provable*. Oxford and New York: Clarendon Press, 1977.
- Mark Colyvan, Helen M. Regan, and Scott Ferson. Is it a crime to belong to a reference class? *Journal of Political Philosophy*, 9(2):168–181, 2001. doi: 10.1111/1467-9760.00123.
- E. Conee and R. Feldman. The generality problem for reliabilism. *Philosophical Studies*, 89(1): 1–29, 1998. doi: 10.1023/A:1004243308503.
- Jonathan Dancy. *An Introduction to Contemporary Epistemology*. Blackwell, 1985.
- Philip A. Ebert, Martin Smith, and Ian Durbach. Lottery judgments: A philosophical and experimental study. *Philosophical Psychology*, 31(1):110–138, 2018.
- David Enoch and Talia Fisher. Sense and "sensitivity": Epistemic and instrumental approaches to statistical evidence. *Stanford law review*, 67(3):557–611, 2015.
- David Enoch, Levi Spectre, and Talia Fisher. Statistical evidence, sensitivity, and the legal value of knowledge. *Philosophy and Public Affairs*, 40(3):197–224, 2012.
- David Enoch, Talia Fisher, and Levi Spectre. Does legal epistemology rest on a mistake? on fetishism, two-tier system design, and conscientious fact-finding. *Philosophical Issues*, 31(1): 85–103, 2021.
- Giada Fratantonio. Evidence, risk, and proof paradoxes: Pessimism about the epistemic project. *International Journal of Evidence and Proof*, 2021.
- Georgi Gardiner. Legal epistemology. In *Oxford Bibliographies Online*. 2019.
- Georgi Gardiner. Profiling and proof: Are statistics safe? *Philosophy (London)*, 95(2):161–183, 2020.
- Georgi Gardiner. Legal burdens of proof and statistical evidence. In James Chase and David Coady, editors, *The Routledge Handbook of Applied Epistemology*. Routledge, forthcoming.

- Sandy Goldberg. Epistemic entitlement and luck. *Philosophy and Phenomenological Research*, 91(2):273–302, 2015. doi: 10.1111/phpr.12083.
- Alvin I. Goldman. A causal theory of knowing. *Journal of Philosophy*, 64(12):357–372, 1967. doi: 10.2307/2024268.
- John Greco. Worries about pritchard’s safety. *Synthese*, 158(3):299–302, 2007. doi: 10.1007/s11229-006-9040-1.
- John Greco. *Better safe than sensitive*, page 193–206. Cambridge University Press, 2012. doi: 10.1017/CBO9780511783630.015.
- Susan Haack. *Evidence Matters: Science, Proof, and Truth in the Law*. Cambridge University Press, 2014.
- Brian Hedden and Mark Colyvan. Legal probabilism: A qualified defence. *Journal of Political Philosophy*, 27(4):448–468, 2019. doi: 10.1111/jopp.12180.
- Avram Hiller and Ram Neta. Safety and epistemic luck. *Synthese*, 158(3):303–313, 2007. doi: 10.1007/s11229-006-9041-0.
- Christopher Hitchcock. Probabilistic Causation. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Spring 2021 edition, 2021.
- H. L. Ho. *A Philosophy of Evidence Law: Justice in the Search for Truth*. Oxford University Press, 2008.
- Jonathan Jenkins Ichikawa and Matthias Steup. The Analysis of Knowledge. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2018 edition, 2018.
- Daniel Kahneman, Paul Slovic, and Amos Tversky. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press, 1982.
- Christoph Kelp. Knowledge: The safe-apt view. *Australasian Journal of Philosophy*, 91(2): 265–278, 2013. doi: 10.1080/00048402.2012.673726.
- Christoph Kelp. Epistemic frankfurt cases revisited. *American Philosophical Quarterly*, 53(1): 27–37, 2016.

- Jennifer Lackey. Epistemic luck [review]. *Philosophical Quarterly*, 56(223):284–289, 2006. doi: 10.1111/j.1467-9213.2006.00443.x.
- Clayton Littlejohn. Truth, knowledge, and the standard of proof in criminal law. *Synthese*, 197(12):5253–5286, forthcoming. doi: 10.1007/s11229-017-1608-4.
- Elizabeth F. Loftus. *Eyewitness testimony*. Harvard University Press, Cambridge, Massachusetts, 1996. ISBN 0674287770;9780674287778;.
- Elizabeth F. Loftus. Eyewitness testimony. *Applied cognitive psychology*, 33(4):498–503, 2019.
- Sarah Moss. Pragmatic encroachment and legal proof. *Wiley: Philosophical Issues*, 31(1): 258–279, 2021. doi: 10.1111/phis.12207.
- Sarah Moss. Knowledge and legal proof. *Oxford Studies in Epistemology*, forthcoming.
- Charles R. Nesson. Reasonable doubt and permissive inferences: The value of complexity. *Harvard law review*, 92(6):1187–1225, 1979.
- Jon Newman. Beyond reasonable doubt. *N.Y.U. Law Review*, 68(5):979–1002, 1993.
- Robert Nozick. *Philosophical Explanations*. Harvard University Press, 1981.
- Michael S. Pardo. Safety vs. sensitivity: Possible worlds and the law of evidence. *Legal Theory*, 24(1):50–75, 2018. doi: 10.1017/s1352325218000010.
- Michael S. Pardo. The paradoxes of legal proof: A critical guide. *Boston University law review*, 99(1):233–290, 2019.
- Michael S. Pardo and Ronald J. Allen. Juridical proof and the best explanation. *Law and Philosophy*, 27(3):223–268, 2008. ISSN 01675249, 15730522. URL <http://www.jstor.org/stable/27652649>.
- Nancy Pennington and Reid Hastie. Cognitive and social processes in decision making. In Lauren Resnick, Levine B., M. John, Stephanie Teasley, and D., editors, *Perspectives on Socially Shared Cognition*. American Psychological Association, 1991.
- Richard Posner. *The Economic Analysis of Law*. Brown and Company, 1973.
- Duncan Pritchard. *Epistemic Luck*. Oxford University Press UK, 2005.

- Duncan Pritchard. Anti-luck virtue epistemology. *Journal of Philosophy*, 109(3):247–279, 2012. doi: 10.5840/jphil201210939.
- Duncan Pritchard. Knowledge cannot be lucky. In Matthias Steup and John Turri, editors, *Contemporary Debates in Epistemology*. Blackwell, 2013.
- Duncan Pritchard. Anti-luck epistemology and the gettier problem. *Philosophical Studies*, 172(1):93–111, 2015. doi: 10.1007/s11098-014-0374-0.
- Duncan Pritchard. Legal risk, legal evidence and the arithmetic of criminal justice. *Jurisprudence*, 9(1):108–119, 2018.
- Mike Redmayne. Exploring the proof paradoxes. *Legal Theory*, 14(4):281–309, 2008. doi: 10.1017/s1352325208080117.
- Lewis Ross. Justice in epistemic gaps: The proof paradox revisited. *Philosophical Issues*, 31(1): 315–333, 2021a. doi: 10.1111/phis.12193.
- Lewis Ross. Rehabilitating statistical evidence. *Philosophy and Phenomenological Research*, 102(1):3–23, 2021b.
- Lewis D. Ross. Legal proof and statistical conjunctions. *Philosophical Studies*, 178(6):2021–2041, 2020a. doi: 10.1007/s11098-020-01521-z.
- Lewis D. Ross. Recent work on the proof paradox. *Philosophy Compass*, 15(6), 2020b. doi: 10.1111/phc3.12667.
- Andrea Roth. Safety in numbers? deciding when dna alone is enough to convict. *New York University law review (1950)*, 85(4):1130, 2010.
- Martin Smith. *Between Probability and Certainty: What Justifies Belief*. Oxford University Press UK, 2016.
- Martin Smith. When does evidence suffice for conviction? *Mind*, 127(508):1193–1218, 2018. doi: 10.1093/mind/fzx026.
- Ernest Sosa. How to defeat opposition to moore. *Philosophical Perspectives*, 13:137–49, 1999. doi: 10.1111/0029-4624.33.s13.7.
- Ernest Sosa. *Judgement and Agency*. Oxford: Oxford University Press, 2015.

- Kathleen Steele and Mark Colyvan. Meta-uncertainty and the proof paradoxes. *Unpublished Manuscript*, 2022.
- Judith Jarvis Thomson. Liability and individualized evidence. *Law and Contemporary Problems*, 49(3):199–219, 1986. ISSN 00239186. URL <http://www.jstor.org/stable/1191633>.
- Laurence H. Tribe. Trial by mathematics: Precision and ritual in the legal process. *Harvard Law Review*, 84(6):1329–1393, 1971. ISSN 0017811X. URL <http://www.jstor.org/stable/1339610>.
- Peter Unger. An analysis of factual knowledge. *Journal of Philosophy*, 65(6):157–170, 1968. doi: 10.2307/2024203.
- Rafal Urbaniak and Marcello Di Bello. Legal Probabilism. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2021 edition, 2021.
- Timothy Williamson. *Knowledge and its Limits*. Oxford University Press, 2000.
- Timothy Williamson. Reply to cassam. In Duncan Pritchard and Patrick Greenough, editors, *Williamson on Knowledge*. Oxford University Press, 2009.
- Linda Zagzebski. The inescapability of gettier problems. *Philosophical Quarterly*, 44(174): 65–73, 1994. doi: 10.2307/2220147.