ABSTRACT

Title of dissertation: A New Theory of Individualized Evidence

Charles Arthur Barclay,
Doctor of Philosophy, 2021

Dissertation directed by: Professor John F. Horty
Department of Philosophy

Do courts treat all types of evidence the same? To the average person, the answer may seem to be *yes*. It turns out, however, that there are some forms of evidence that the law prefers over others. To see this, consider a pair of cases that are, for the most part, very similar:

Case 1[1]

[1] This case is based on Smith v. Rapid Transit (1945), where claimant was struck by a hit-and-run bus and based her claim that the bus was the defendant's solely on evidence that the defendant operated 90 percent of the buses in the city. The Massachusetts Supreme Judicial Court sustained defendant's motion for summary judgment on grounds that that the base rate statistic was insufficient to make a case against the defendant in the absence of more particularized proof of the ownership of the offending bus.

Suppose it is late at night, and claimant's car is hit by a black car from one of two ride-sharing companies in the area, A and B. Claimant cannot identify whether the car that hit her belongs to A or B, but she can prove that 85 percent of the black, ride-sharing cars in the city are operated by A, and that the remaining 15 percent are operatd by B. Moreover, each of the other elements of the case - negligence, causation, and especially, the fact and the extent of the injury - is either stipulated or established to a virtual certainty.

Case 2

Again, claimant's car is hit by a black car from one of two ride- sharing companies in the area, A and B. But, unlike in case 1, each company operates 50 percent of the ride-sharing cars. Furthermore, an eyewitness establishes that the black car belongs to company B. Assume that the witness is usually reliable during the day; however, the incident occurs at night when she is prone to make a mistake 25 percent of the time. Finally, assume that each of the other elements of the case - negligence, causation, and especially, the fact and the extent of the injury - is either stipulated or established to a virtual certainty.

Given the facts of each case, our confidence that a vehicle from Company A hit claimant in the first case is 0.85; but our confidence that a veheicle from Company B hit claimaint in the second case is *only* 0.75.[2]

[2] Some might argue that our credence in the second case is not actually 0.75. On this view, credences ought to take into account the total available evidence and since each case presents a piece of evidence, the claimant's credence in case 2 that company B is at fault should reflect the evidence each case offers. But this is orthogonal to the key question that is addressed in this paper. Thus, I shall treat the two cases separately. This means that the credence produced in case 2 will not depend on any information available in case 1, and vice versa.

Despite the probative value of the first case being higher than that in the second, in a set of experiments where subjects were given a pair of similar cases, Gary Wells (1992) shows that many people, judges and lay people alike, are open to finding for the claimant in the second case, but *not* in the first. The unequal treatement of both cases even extends to the law. Most courts would throw out the first case while letting the other case go to trial. The main reason for the different treatment of each case centers on the type of evidence that each claimant appeals to. To be sure, the claimant in the second case relies on eyewitness testimonial evidence. This type of evidence is an instance of individualized evidence - *viz.,* evidence that is specific to the individual. But, the claimant in the first case rests solely on bare (or, naked) statistical evidence.

Discussions of what marks the distinction between individual and bare statistical evidence has garnered some attention in epistemology. Proposals abound in explaining what makes one piece of evidence individualized and not another. In what follows, I show that several prominent frameworks in the literature do not have desirable properties that we might want in a theory of individualized evidence. I then defend a novel approach based primarily on a relevant alternatives account of knowledge view in epistemology. The view, as I will show, does possess several properties that are significant for a theory of individualized evidence.

A NEW THEORY OF INDIVIDUALIZED EVIDENCE

by

Charles Barclay

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2021

Advisory Committee:

Professor John F. Horty, Chair
Professor Marcello DiBello (ASU)
Professor Eric Pacuit (UMD)
Professor Christopher Morris (UMD)
Professor Piotr T. Swistak (UMD)

I dedicate this thesis to my family and loved ones for nursing me with affections and support.

ACKNOWLEDGMENTS

Six years ago, I signed up for a graduate seminar on reasoning in the law, which was taught by Jeff Horty. At the time, my goal was modest: I aimed to learn about the applications of epistemology and logic to evidence law theory. I had no idea that the class would spark a deep-seated interest in me that would shape the direction in which I would go in my dissertation. In more ways than one, the course was deeply enriching. I wrote a paper that eventually led to a publication. I also learned of proof theories that have been developed to model the type of reasoning that lawyers and legal experts do. When Jeff noticed my interest in the literature, he began to cultivate my knowledge through non-intrusive means, which ranged from placing journal articles and books on reasoning in the law in my mailbox in the Department to sending me links to papers and course syllabi that had a reading list of journals that I should consult. Much of my knowledge in the literature of philosophy of law was initially shaped and directed by Jeff.

By the Spring of my second year, I took, what would turn out to be, my final seminar at UMD. The course was a grad-level seminar on Belief Revision - or at, least that's what I thought it was. Eric Pacuit taught the course. Of course, out of all the committee members of my dissertation, I've known Eric the longest. Prior to coming to UMD, he had spent a semester in the Philosophy Department at CMU, which overlapped with when I was at the tail-end of my Master's program in the same department. He and I met once a week to discuss possible worlds. We began reading Lewis' and Stalnaker's possible world semantics. By the end of the semester, I had a firmer grasp on the literature of possible worlds. But more importantly, I had

a newfound appreciation for Math and its connection to Philosophy, thanks to Eric. Those were some of the memories I took with me when I signed up to take Eric's class. Little did I know that I would read a paper that would ultimately lead to the topic that my dissertation covers. The paper was *Belief, Credence, and norms* by Lara Buchak. In one of her arguments, Buchak asserts that naked statistical evidence cases should not be treated the same as eyewitness testimonial cases. I thought her view was obviously false. To prove Buchak wrong, I resolved to write a paper that supported the plausibility of naked statistical evidence in some cases. After some reflection, I titled the paper *Is Naked Statistical Evidence Less Secure Than Eyewitness Testimonial Evidence*?

In the ensuing semester, Aleks Knocks, my then officemate and fellow graduate student, informed me of the ICAIL Workshop on 'Evidence and Decision Making in The Law", which would be held in June in London at King's College. The paper that I had written in Eric's seminar was selected for a presentation at the workshop. Despite being allotted a short time to present my work, my time at the workshop was worthwhile. Not only did I receive great feedback on my work, but I also learned a range of positions on naked statistical evidence. But, as good as my time in London was, the greater benefits of attending the workshop were still to come. I forged relationships with several attendees of the workshop, one of whom is Henry Prakken - professor in Informatics and legal Argumentation at the Law Faculty of the University of Groningen. Henry provided some feedback on the paper that I presented at the conference. Moreover, he gave invaluable suggestions and insight on another paper, which led to a publication. When it was time to select committee members on my dissertation, Jeff and Eric agreed to be on the committee. Henry

also followed suit.

Another person with whom I stayed in touch during my visit to King's College is Marcello DiBello, one of the organizers of the workshop. After discussing it over with Eric, I also asked Marcello to be on my committee. He acquiesced and let me know that he was sympathetic to my position at the time. Marcello's involvement in my dissertation was nothing short of impactful. On several papers that I wrote, some of which I had no intention of including in my dissertation, he provided copious feedback in a timely manner. Marcello was also liberal with his time. On many occassions, he met with me over skype or zoom to provide specific guidance and support. If that wasn't enough, he even recommonded other types of research that I could pursue that related to my dissertation. I will always be grateful for his constructive feedback and the high standard that he held me to.

In late September of 2019, my dissertation reached a key turning point. I met with Jeff over lunch to discuss where I was in my research. At the time, I was planning to submit three separate but related, publishable papers for my dissertation. One had already been written, and another was in the process of being written. After sharing with Jeff my interest in naked statistics and the range of positions in the literature, he then suggested that I consider developing a theory that states when it is appropriate to use naked statistical evidence. The suggestion resonated with me. Given my interest in the topic, it seemed sensible to write a paper on naked statistics. And, what better paper to write than one that shows in which cases it may be reasonable to find for the plaintiff who relies on bare statistics? The trouble, however, was coming up with a plausible theory that not only explains why naked statistics are problematic, but also shows in which cases it may be permissible to use

them.

My doubts proved correct. Coming up with a robust theory was not easy. Many of my ideas hit a dead end. That all changed, however, one cold afternoon in November, as I was getting ready to teach an evening class. While thinking about the ways in which epistemology can aide and inform issues in the law, it suddenly hit me: the problem with naked statistical evidence can be modeled by a relevant alternatives account of knowledge in contextualism. That is, in the contextualism literature, one popular view of knowledge states that to have knowledge of some empirical proposition p just means that the agent has eliminated all relevant alternatives to p. Similarly, when the trier of fact finds for the plaintiff, she has done something akin to eliminating relevant alternatives to the plaintiff's narrative.

Although my idea was in its embryonic stage, I ran it by some committee members and soon realized that my view needed some work. I subsequently wrote a draft in which I argued that, with a few alterations, a relevant alternatives account of knowledge from the contextualism literature has the resources to model what is deficient about naked statistical evidence. Through the help of two committee members of my dissertation, I learned that Sarah Moss had already written a similar paper in "Knowledge and Legal Proof". For obvious reasons, it was heartbreaking to learn that someone had beaten me to my idea. The one consolation, however, was seeing that there might be something to my idea after all, if a well-respected philosopher won an award based on virtually the same idea. Despite the setback, I decided not to abandon my core idea of using a relevant alternatives account of knowledge to model a theory of evidence. Instead, I looked for ways in which my view might differ from Moss'.

After a series of discussions with some members of my committee, and with Jeff's prodding, I realized that I had enough to write a comprehensive paper that would constitute the entire dissertation, rather than just letting that paper serve as one of three papers making up the dissertation. I believe that I have a theory of individualized evidence that is sufficiently different from that of Moss'. Of course, my efforts could not have been reached without the support, insight and cajoling of my committee. When I think of their collective support and guidance, I am reminded of a piece of advice that Zvi Bienner, one of my former professors, gave me. After I asked him how I should approach my graduate education, he succinctly replied: "Think of graduate school as on-the-job training." Through the efforts of my committee members, I believe that I have been trained as a philosopher, college professor, researcher, and colleague!

TABLE OF CONTENTS

INTRODUCTION

This dissertation attempts to make several contributions to the current debate in the philosophy of law literature on what constitutes individualized evidence - *viz.,* evidence that is specific to the individual. The main question has evolved throughout time. Initially, the question was put this way:

(A) Why do courts tend to disfavor cases in which the plaintiff relies solely on statistical evidence (sometimes called naked statistical evidence cases) but not those that rely solely on eye-witness testimonial (or, other forms of individualized) evidence?

For convenience, some refer to *(A)* as the naked statistical evidence problem. Eventually, the above question morphed into a question of what makes statistical evidence different from individualized evidence. Or, simply put:

(B) What is individualized evidence?

> To this second question, people typically respond in different ways:
>
> Some view the question is best answered by appealing to ideas and concepts from epistemology. Others view the question can be answered by stating the moral and legal concerns with naked statistical evidence cases that individualized evidence cases are able to avoid.
>
> And, still others believe that some combination of the first two answers is where the correct answer lies.

Each response has its own merits. But, the plurality of answers to *(B)* come at

several costs. One worry is as follows: By allowing an array of answers with little to no guidelines, it allows authors to speak past each other without much progress being made. To curb this problem, the first contribution I make is to offer four desirable properties that a robust theory of individualized evidence ought to possess. By proposing criteria, I hope to shape the debate and offer a new way in which discussions on individualized evidence may proceed.

The desirable properties I propose range from requiring theories to explain the exceptional treatment of DNA evidence in the law to requiring that the theories must adapt in accordance with the different standards that the case is in. There is even a requirement that imposes a restriction on how frameworks satisfy some of the other desirable properties. Chapter 2 states the desirable properties in fuller detail.

The reader will be curious to know my own answer to *(B)*. En route to providing my answer, I show just how challenging the second question is by stating the numerous answers scholars have given to the question. As David Enoch and Talia Fisher note, "A good way to appreciate the depth of a problem is to explore the attempts that have been made to tackle it." (Enoch and Fisher; 2015, 565). My discussion of historical answers to *(B)* is provided in Chapter 3.

My answer to the second question constitutes my second contribution to the literature. I develop a theory of individualized evidence, which I call Legally Relevant Alternatives (or, *LRA*). *LRA* is partly insipred by a recent debate in epistemology called contextualism, and is fully explained in Chapters 4 and 5. I conclude my dissertation by discussing the other applications of *LRA*. For example, *LRA* provides a highly plausible response to the *Gatecrasher Paradox*.

Here is a summary of what is to come:

Chapter 1 - "The Problem with Statistical Evidence" - introduces the problem with using bare statistics to base a verdict. I discuss two well-known examples that illustrate the problem. I briefly mention the myriad of proposed solutions to the problem.

Chapter 2 - "Desirable Properties for Individualized Evidence" - proposes four highly plausible desiderata that a robust theory of individualized evidence should possess. Each criterion is best understood as a necessary condition for individualized evidence. I do not claim to have originally come up with the first three properties. Indeed, several philosophers of law and epistemologists have discussed these ideas or precursors to them. The fourth criterion is my own original contribution to the literature as it relates to desirable properties for a theory to possess.

Chapter 3 - "Current Theories" - reviews the literature on individualized evidence. This literature has had many contributions from legal scholars, philosophers of law and epistemologists. The contributors explain what makes individualized evidence different from statistical evidence. In most cases, their explanations amount to a full-blown framework of individualized evidence. I explain the frameworks and assess whether they are able to satisfy the desiderata stated in Chapter 2.

Chapter 4 - "Epistemology" - marks the start of my positive proposal. I discuss the background to my theory by explaining a fairly recent movement in epistetmology called contextualism. I then explain the relevant alternatives account of knowledge view (hereafter, *RAA*), which is a popular contextualist approach in epistemology. Ultimately, the discussion in this chapter shows the relationship that my proposal

has with epistemology. It also will make my proposal easier to understand.

Chapter 5 - "From Epistemology to the Law" - unveils my theory of individualized evidence: Legally Relevant Alternatives (or, *LRA*). I also explain in which ways *LRA* from two other theories. The first is *RAA*, which is the contextualist theory in epistemological on which my proposal is based. The other is Moss' probabilistic account of knowledge, which also borrows concepts and ideas from *RAA*.

Chapter 6 - "LRA and the Desiderata" - attempts to set *LRA* apart from the contemporary frameworks stated in Chapter 3. I show that *LRA* respects all four of the desirable properties stated in Chapter 2. I also deal wth several possible objections to my account.

Chapter 7 - "LRA and the Gatecrasher - introduces a legal puzzle based on statistical evidence first introduced by Laurence J. Cohen called the *Gatecrasher Paradox*. I discuss several typical responses legal scholars have had to the *Gatecrasher* before explaining *LRA*'s own solution. I conclude by showing how *LRA* can captue some of the historical responses to the *Gatecrasher.*

CHAPTER 1

THE PROBLEM WITH NAKED STATISTICS

Quantitative evidence has played a role in court decisions. In the sentencing phase, for example, a judge may use statistics to determine the impact that a lighter sentence may have on recidivism. The trial phase also admits quantitative information. The plaintiff may use statistics as supplementary evidence to establish a link between some harm she experienced and the defendant's misconduct.[3]

But, the acceptance of quantitative evidence is not without limitations. Consider, for example, naked statistical evidence cases, where the only evidence the plaintiff offers is statistical. In the first half of the twentieth century, where such cases began to appear in court, many judges deemed it problematic to find for the plaintiff who only relies on bare statistics. But the doctrinal picture on statistical evidence is a bit more complex than this.[4] Not only have courts deemed statistics irrelevant in many cases, but in other cases statistical evidence has sufficed for a finding for the plaintiff. In *Turner v. United States* (1970), for example, the Supreme Court held that base-rate evidence indicating that more than ninety-nine percent of all heroin consumed in the United States is illegally imported was enough to allow an inference that its possession amounted to possession of a smuggled drug.

[3] In this essay, I will refer to plaintiff and claimant interchangeably.

[4] (See Fisher et al., 2015, 561)

Unfortunately, discussions of naked statistical evidence among legal scholars aren't less divisive. Although most evidence scholars believe that it is unacceptable to find for the plaintiff in such cases, there is much disagreement on why it is unacceptable. One reason for the disagreement can be attributed to their methodology. Scholars tend to point to different characteristics of naked statistical evidence as the root problem. Such problems are traditionally captured by several examples. Given the role that these examples will serve throughout the paper, I will discuss two of them now: the *Prison Riot Case* and the *Blue Bus Case*.[5] It is important to note that these examples are imaginary cases, despite being based on actual cases. I will refer to them as 'cases', but this does not imply that they are real, or realistic.

Charles Nesson (1979) introduced the *Prison Riot Case*. I will use the following version.

100 prison inmates participate in a riot and escape from the prison. During the riot, 99 of them assault and kill a guard on duty. One prisoner is caught, and is prosecuted for participating in the killing of the guard. It is absolutely certain that he is one of the 100 prisoners that participated in the riot, and it is absolutely certain that 99 of them participated in the killing of the guard. There is no further evidence (no witnesses, no CCTV images) that makes it possible to identify the prisoners who participated in the killing, and individuate the prisoner who did not participate. According to the prosecutor, the evidence that the defendant is one of the prisoners who participated in the riot is sufficient for the standard of proof in criminal cases, since it establishes that there is a 99% probability that the defendant participated in the killing of the guard. (If someone would object that the standard of

[5]A third example is the *Gatecrasher Case*, which will be discussed in Chapter 7.

proof in criminal cases corresponds to a probability higher than 99%, say 99.5%, the circumstances of the case can always be altered to satisfy this number, for example by making it 1000 prisoners where 999 participated in the killing of the guard.)

Nesson maintains that the defendant should not be convicted on such evidence. The plaintiff does not provide evidence that singles out the defendant from other prisoners.[6] It is therefore wrong to find for the plaintiff in this case and all cases involving naked statistical evidence.[7]

Finally, Judith Jarvis Thomson (1986) introduced the *Blue Bus Case*. One version of it runs as follows:

On a rainy night a pedestrian is hit by a bus and injured. It is certain that the vehicle was a bus, and there is no doubt that the driver was negligent. Several buses were in traffic on the night in question, and it cannot be established which of them hit the pedestrian. Two bus companies—the Blue Bus Company and the Green Bus Company—operate in the area, so the bus that hit the pedestrian belongs to one of these companies. Both companies say that they don't know if the accident was caused by one of their buses, or by the other company. The Blue Bus Company has the major market share. 80% of the buses who could have hit the pedestrian belong to the Blue Bus Company, and only 20% belong to the Green Bus Company. The pedestrian therefore sues the Blue Bus Company. There is an 80% probability that the damage was caused by the Blue Bus Company, and the standard of proof

[6] (See Nesson; 1979, 1193)

[7] (See Nesson, 1985, 1378)

in tort cases only requires that the plaintiff makes his case more probable than the defendant's case ('preponderance of the evidence').

Similar to Nesson, Thomson argues that it would be wrong to judge the plaintiff on the basis of naked statistics. To impose liability on the defendant requires a different kind of evidence that Thomson describes as 'individualized evidence'[8], as in a witness that observes identifying information of the bus that hit the pedestrian. Naked statistical does not differentiate the culprit in this way, and is therefore not sufficient for the burden of proof.

Recently philosophers have weighed in on the problem of naked statistical evidence. Many of them view the problem as epistemic in nature. The contributions made by philosophers and legal scholars have led to a wide varierty of different solutions to the problem of naked statistical evidence. It has been suggested that it is a problem of *social acceptability* (Charles Nesson), *individual autonomy* (Wasserman, Pundik) *being a member of the wrong reference class* (Colyvan et al.), *lack of a causal connection to the facts* (Thomson), *sensitivity with regard to incentives* (Enoch, Spectre & Fisher), or *normalcy of explanation* (Smith).

From the various proposals, the following questions emerge: What is the best way to evaluate the different proposals? And, is there a better theory that has not yet been proposed? This paper aims to answer both questions. My main goal is to offer a theory of individualized evidence that possesses several highly plausible properties, which cannot all be found in any of the abovementioned proposed solutions. In order to show the strengths of my theory, it is first necessary to discuss several desiderata

[8] (See Thomson; 1986, 203)

that a robust theory of individualized evidence ought to possess. The next chapter aims to do that.

CHAPTER 2

Desirable Properties for Individualized Evidence

2.1 THE NEED FOR DESIDERATA

The debate on what constitutes individualized evidence is plagued by several issues. One issue seems to be a matter of ambiguity. Simply put, the debate is not well defined. This concern can easily be seen by a brief look at the various ways in which the debate has been framed throughout the years. Enoch and Fisher (2015) see the debate as a question of why courts exhibit a general preference for individualized evidence over statistical evidence for imposing liability. This then motivates them to provide an overarching theory of individualized evidence.[9] Moss (2021) sees the debate as a question of what legal proof requires. By framing the debate in that way, she then proposes an account of legal proof that adapts to each legal proceeding.[10] For other contributors, the debate is really a question of moral acceptability. For Colyvan et al., (2001), for example, the debate centers on whether the amount of illegal substance found on the date of a convicted felon's arrest should be used as an estimation of previous acts of wrong doing.[11] Colyvan and colleagues then argue

[9] (See Enoch and Fisher; 2015, 564)

[10] (See Moss; 2021, 4)

[11] Specifically, their concern is about whether prompted by an appellate court va-cating a lower court's decision. The lower court based sentencing for a convicted.

that it should not be a crime for belonging to a reference class.[12]

It is important to note that the common thread in each above-mentioned article concerns the shortcomings of naked statistics. However, it is also clear that the debate has meant something slightly different to each contributor. We should not think that this is a problem we can afford to ignore. Not having a well-defined debate can easily lead to authors talking past each other, which in turn can make it difficult to assess a positive contribution.

Another issue with the debate on individualized evidence is an issue of guidelines. Specifically, the debate lacks clearly articulated guidelines or criteria that proposed solutions must satisfy. Indeed, this issue is related to the previous one. If the debate is a moving target, then guidelines for assessing positive contributions will also be hard to set. While the literature is rife with proposed solutions to the debate, there is no consensus on a set of criteria that a robust theory ought to possess.

The issues I have identified so far have created, among other things, a need for a set of criteria that a theory of individualized evidence ought to display. The objective of this chapter is to propose four highly plausible properties that we should look for in a robust theory of individualized evidence. With the exception of the last property, I do not claim to have originally come up with the first three properties. Indeed, the first three properties have been mentioned in several articles in the literature.

felon on an estimation of the total quantity of drugs he is believed to have smuggled into the USA.

[12](See Colyvan et al., 2001, 168-170)

2.2 DNA EVIDENCE

If a theory of individualized evidence should satisfy several desiderata, then what are they? More importantly, what makes them good properties to possess. Given the issues with naked statistics laid out in the first chapter, we might wonder whether the problems apply to all forms of statistical evidence irrespective of how the data are generated. When looking at the history of legal decisions involving statistical evidence since the mid 1980s, a curious phenomenon stands out: Despite the general approach of inadmissibility of naked statistics, the courts are willing to make an exception when it comes to a particular type of statistical evidence: DNA evidence.[13]

The exceptional treatment in the law that DNA evidence enjoys cannot be overstated. But how is this possible, given the court's hostitility towards other types of naked statistics? What makes DNA evidence different from other forms of statistical evidence? If a theory of individualized evidence is to reflect how courts handle most cases of individualized evidence, then it seems reasonable to require it to account for the exceptional treatment the law renders to DNA evidence. My first desideratum then requires that frameworks of individualized evidence properly account for the exceptional treatment of DNA statistical evidence.

Alas, there is a challenge to this desideratum. Many people are unaware that DNA evidence is a type of statistical evidence. Some may even deny the relation between DNA evidence and statistics. To motivate the first desirable property, the next Subsection shows why DNA evidence is seen as a type of statistical evidence. In the

[13] (See Roth; 2010)

Subsection thereafter (i.e., Subsection 2.2.2), I elaborate on the courts' acceptance of DNA evidence.

2.2.1 Why DNA evidence is a form of statistical evidence

One context in which DNA evidence is used is, the so-called, 'cold-hit'[14] DNA evidence. In such a context, the genetic material found at the crime scene is compared to offenders residing in the Combined DNA Index System (CODIS) - a computer network that connects forensic DNA laboratories at the local, State, and national levels.[15] CODIS databases contain DNA profiles of two general kinds: known profiles and crime-scene profiles. The sources of the known profiles include voluntary submissions and mandatory contributions from certain convicted offenders and in some cases arrestees.[16] Crime-scene profiles are collected by crime-scene technicians and are evidentiary samples connected to unsolved offenses according to certain standards related both to the quality of the evidence and the clarity of its connection to

[14] Cold hits raise legal questions. Given that it is possible to recover biological material from well-preserved evidence, many cold-hit cases have arisen from offenses that occurred years or even decades ago. Unfortunately, clear data on the number, frequency, and outcome of cold-hit cases are difficult to obtain. Some information exists, however. In Virginia, a survey of the outcome of the first 1000 cold hits revealed that 100 resulted in convictions through plea or trial, seven yielded not guilty verdicts, and 53 were never prosecuted; 752 were pending at the time of the survey.

[15] *Ibid*

[16] (See Butler; 2005)

the crime.

In many cases, there is a match between the DNA sample taken at the crime scene and a DNA profile in CODIS. The match is then used as the basis of a criminal case against the person to whom the matching DNA profile belongs. Although it is unlikely that a criminal suspect would share a DNA profile with a piece of incriminating evidence by coincidence, quantifying this probability requires estimating the population frequency of the varying genetic features (alleles) in a specified reference class, such as a racial or ethnic group.[17]

In the simplest computation, the population frequencies of each of the varying genetic features are deemed to be inherited independently of one another. They are then multiplied together, along with various coefficients, in keeping with the product rule.[18] The result is commonly known as the "random match probability," and it is equivalent to the chance of a DNA match by coincidence. It is also the statistic usually reported to a judge or jury. This statistic indicates the estimated rarity of the DNA profile, although that specific figure may be misleading if it induces the fact-finder to ignore the much more likely chance of a laboratory error or errors in subjective judgments made in interpreting the results.[19]

[17] (See, Lindsey et al., 2003)

[18] *Ibid*

[19] One component of DNA evidence is extracting, analyzing, and making statistical inferences from genetic material. Suppose, for example, the perpetrator left a cup containing saliva at the scene of the crime. This genetic material is processed in a laboratory to build a profile of which alleles appear at which locations on a strand of DNA. From this process, DNA evidence is used in different contexts. One is where suspicion has already fallen upon some person and DNA profiliing is used to bolster against them. This often involves taking further genetic material from a suspect and

2.2.2 The Exceptional Treatment of DNA evidence in the law

The number and frequency of cold hits has grown rapidly. In November 2004, the FBI reported a total of 19,500 cold hits, including both scene-to-scene and offender-to-scene matches[20] by March 2007 that number had risen to over 47,000. States have experienced similar growth. For example, it took the state of Virginia from 1993 to 2001 to reach its first 1000 cold hits, but the second thousand occurred within the following 18 months.[21]

Cold-hit cases have prompted courts to confront the question of whether a genetic match constitutes sufficient evidence to uphold a conviction. By comparison, the Supreme Court has previously ruled that a confession, standing alone, cannot serve as the basis for conviction.[22]Although it has not yet been fully decided, most courts have resolved the cold-hit question in the affirmative. As one court observed, "the perils of eyewitness identification testimony far exceed those presented by DNA expert testimony"[23] And another, while recognizing that DNA evidence is not "infallible,"

testing it against the initial genetic material found at a crime scene. If the samples match, then this further supports the pre-existing suspicion.

[20] NDIS Statistics, FBI, Measuring Success, 2005. Available at http://www.fbi.gov.proxy-um.researchport.umd.edu/hq/lab/codis/success

[21] Murphy E. The new forensics: criminal justice, false certainty, and the second generation of scientific evidence, 56 Cal. L. Rev. 721, 2007:742

[22] (See *Smith v. United States*, 348 U.S. 147, 152 (1954))

[23] (See *Roberson v. State,* 16 S.W.3d 156, 170 (Tex. Crim. App. 2000))

explained that "[v]irtually no evidence is absolutely conclusive"[24] Of course, cold-hit cases raise justice-related concerns, especially since mounting a defense to a crime that occurred in the past becomes increasingly difficult as time progresses.

Cold-hit cases can vary strikingly in terms of the evidentiary value of the DNA match itself. This is especially true depending on whether there is corroborating evidence. For example, such evidence may be weak, as in the information that the defendant lived near the victim.[25] The evidence could also be strong, as in the defendant matches a detailed physical description given by the victim at the time of the offense. In *People v. Johnson*, a 15-year-old rape victim provided descriptions of the assailant's car and physical features, including distinctive tattoos, but no suspect was identified. Five years later, a cold hit identified the defendant, and the ensuing investigation revealed that he had lived in the area at the time of the offense, had matching tattoos, and owned a vehicle that fit the description given by the victim.

Courts across the country have also upheld convictions based only upon cold-hit evidence. In cases of sexual assault, courts have reasoned that the intimate nature of the sample decisively puts an end to arguments that it might have been left inadvertently. For example, in *State v. Hunter*, an appellate court upheld a rape conviction based on a semen sample collected after a sexual assault in 1995. The sample was matched eight years later to the defendant, although literally no other evidence linked him to the crime.[26] In sustaining a rape conviction based on DNA

[24] (See *People v. Johnson* 139 Cal.App. 4th 1135 Cal. Ct. App. 2006).

[25] (See *Riggs v. State*, 809 N.E.2d 322 Ind. 2004)

[26] (*State v. Hunter*, 861 N.E.2d 898 (Ohio App. Ct. 2006)).

evidence alone, another court explained, "we cannot conceive of an innocent reason for the defendant's DNA to be found on swabs taken from the victim's anal area."[27]

Some courts have upheld prosecutions based on cold hits not only in the absence of additional evidence but also in the face of contrary evidence. For instance, in Michigan, the government tested a biological sample collected from the hand of the victim in a 36-year-old murder case. The test revealed two profiles: one of the man ultimately prosecuted for the offense and another of an individual who was 4 years old at the time of the murder. Despite the unexplained presence of the second profile and the absence of any additional evidence, the jury convicted the defendant.[28] Likewise, in *United States v. Jenkins*, the court approved a murder and burglary prosecution. The conviction was based primarily on genetic evidence, despite the fact that a day after the incident another man was found in possession of the decedent's credit cards and other items taken from the home.[29] Although the prosecution was allowed to proceed, the case resulted in a mistrial.[30]

In sum, despite the statistical nature of DNA evidence, the history of its usage in

---

[27] *Ibid.* At the same time, however, the court specifically disclaimed "an iron-clad principal [sic] that DNA evidence, without corroboration, is always sufficient," since "[p]ractically infinite factual variations can arise."

[28] (See Convicted Murderer Seeks Retrial, Kalamazoo Gazette, May 10, 2006; Sect. Local News)

[29] (See United States v. Jenkins)

[30] (See Cauvin HE)

legal settings cannot be overstated. Ever since its introduction in American courtrooms in the 1980s, DNA evidence has proved a powerful tool leading to scores of convictions and post-conviction exonerations.[31] American courts admitted DNA evidence in paternity suits[32] and in criminal trials.[33] Nearly a decade later, most states have allowed DNA test results to be admitted as evidence in criminal trials.[34] In fact, some state legislatures have passed laws recognizing the reliability and admissibility of DNA evidence.

So far a strong case has been made showing most courts' (and even some legislatures') wide acceptance of DNA evidence. But there are questions that this acceptance of DNA evidence raises that must be addressed.[35] One question, for example,

[31] (See *DNA Exonerations Nationwide*)

[32] (See Gilmore; 2002)

[33] The first appellate court in the United States to uphold the admission of DNA evidence in a criminal proceeding was Florida's Fifth District Court of Appeal. (See Andrews v. State). The first federal appellate court to approve of the admission of DNA evidence was the *United States v. Jakobetz*.

[34] (See Whitmore; 1993)

[35] Not everyone believes that DNA evidence cases suce for conviction. Martin Smith, for example, has offered a theory of individualized, according to which such evidence is one that satisfies normic support. He defines normic support as follows: E normically supports p in the sense that if E holds, then p normally holds, and so whenever E holds but p does not, an explanation for the abnormality is required besides the falsity of p. Smith does not think that DNA cold-hit cases satisfy normic

can be put as follows: Why do courts give DNA evidence preferential treatment? In the next Subsection, I argue that one explanation can be found in, what is sometimes called, the Frye test.

2.2.3 Understanding the Frye Test

The Frye test had its origin in *Frye v. United States,* a 1923 United States Court of Appeals decision concerning the admissibility of evidence derived from a systolic blood-pressure deception test - an unsophisticated precursor to the polygraph machine. Alphonzo Frye was convicted of murder in the second degree. In the course of the trial, counsel for the defendant offered an expert witness to testify to the result

support. To see this, suppose in a cold-hit case the defendant who was incriminated by the cold-hit match was someone other than the perpetrator. Furthermore, routine sources of error such as mislabeling of the samples, etc., were ruled out. It could happen, by sheer coincidence, that the perpetrator and the defendant shared the same DNA profile. Smith argues that this would not require normic support. (See Smith; 2018, 1212), because no explanation for why the match coincidentally pointed to the wrong person is required. Given that such cases do not satisfy normic support, this then means that DNA cold-hit cases do not `suffice` for conviction.

I do not think that Smith's argument goes through, however. Di Bello (2020), for example, argues that Smith's argument overgeneralizes to other forms of trial evidence. Thus, while there is some resistance to readily accept DNA statistical evidence as individualizing, it still remains to be seen whether any of the arguments is successful.

of a deception test made upon defendant. The test is better known as the systolic blood pressure deception test.

The basic idea behind the deception test is that changes in blood pressure would be caused by changes in the emotions of the witness, and systolic blood-pressure rises were caused by nervous impulses sent to the autonomic nervous system. Frye argued that scientific experiments confirmed that pain, fear, and rage usually produced an elevation of systolic blood pressure, and that conscious deception or falsehood, concealment of facts, or guilt of crime, accompanied by fear of detection when the person is under examination, raised the systolic blood pressure in a curve, which corresponds exactly to the struggle going on in the subject's mind, between fear and attempted control of that fear, as the examination touches the vital points about which he was attempting to mislead the examiner.

The court held that test results that Frye attempted to introduce into evidence did not meet the requirement that such evidence be "sufficiently established to have gained general acceptance in the particular field in which it belongs," and therefore the test results were properly excluded by the lower court. The court reached its decision on the basis that though the deception test at issue here has a scientific basis, "just when a scientific principle or discovery crosses the line between the experimental and demonstrable stages is difficult to define ... [and] the thing from which the deduction is made must be sufficiently established to have gained general acceptance in the particular field in which it belongs [to be admissible]."[36]

The reasoning behind the court's decision in *Frye* led to, what is now called, the *Frye* test, which asks "whether the accepted techniques, when properly performed,

[36](See *Frye v. United States*)

generate results accepted as reliable within the scientific community generally." The *Frye* test, in some respects, explains the court's acceptance of statistical evidence generated by DNA research. The scientific community has overwhelmingly found this as a reliable method for generating knowledge. The discovery of DNA as well as our understanding of its structure and functioning is arguably the most important discovery of the 19th century. The impact of the discovery of DNA on scientific and medical progress has been massive. Consider, for example, the impact DNA research has had on genetics. Owing to the discovery of DNA, geneticists' ability to actually diagnose diseases early on have vastly improved. Moreover, scientists have been able to better assess a person's genetic susceptibility to specific diseases, which have led for the opportunityto create new drugs and interventions to treat these diseases. Indeed, the scientific commnuty's acceptance of DNA research as a very reliable method for acquiring knowledge has contributed to the court's acceptance of DNA evidence.

2.2.4 The litmus test for the first desideratum

Our discussion of the first desideratum concludes with a way to assess whether a framework respects the requirement. I will use two cases of naked statistical evidence to assess whether the framework satisfies the first desideratum. The first is a cold-hit case and the second is a case the statistics for which are not generated by DNA research:

DNA: Suppose a man is fatally wounded at his home. His death is the result of compression of anatomical neck structures, which lead to asphyxia. After doing a sweep of the crime scene, detectives uncover fingernail scrapings all over the victim's neck. The scrapings enable DNA samples to be extracted and a DNA profile is

generated. A match is then made between the DNA profile recovered from the victim's neck and a DNA databank of offenders. The match supports a 1 in 10^{10} chance that someone other than *Defendant*'s DNA was found on the victim's neck. *Defendant* is subsequently charged with homicide.

Brawl[37]: Suppose twenty one employees of a company have assembled to protest unfair wages. After being told by a police officer that they must disband and leave the premises because of excessive noise, all but one of the protesters attack the officer, who is subsequently injured. *Defendant*, a randomly selected protester, is charged with assaulting an officer.

In each case, there is a greater than .95 chance that *Defendant* actually commits the crime that they are being accused of. However, intuitions vary on whether a guilty verdict should be returned for each case. Most people think that *Defendant* should not be found guilty in ***Brawl*****,** but should be in ***DNA.***[38] To explain the exceptional

[37] This example is obviously inspired by the *Prison Riot Case*. My motivation for using this example instead of the *Prison Riot* is that I want to offer a non-DNA-based case of statistical evidence that, despite generating a high confidence level of being in favor of the plaintiff, would still deemed be unacceptable to return a guilty verdict.

[38] Although I took some effort to carve out, what I believed to be, a convincing example involving DNA evidence, not everyone shares my opinion that we should find for the plaintiff. Some, for example, might wonder why we think that a guilty verdict should be returned in *DNA* just because *Defendant*'s DNA is found on the victim's neck. It seems possible that a number of hypotheses need to be ruled out that don't make it obvious that a guilty verdict should be automatically returned. It could be, for example, that *Defendant* has an identical twin brother who actually

committed the crime. Given that identical twins share their DNA code with each other, this explains why a DNA profile matching that of *Defendant's* was found on the victim's neck. Or, it could be that several hours before the victim was brutally attacked, he was at a physical therapy session and received electrical stimulation treatment on his neck followed by a neck message both of which were provided by *Defendant*. The actual perpetrator who killed victim used gloves and was careful not to leave behind any trace. Each example offers a possible explanation for *Defendant'*s DNA being found on victim's neck, while establishing the likelihood of *Defendant*'s innocence. Given these possibilities, as the objection goes, why do we think it obvious that *Defendant* be found guilty on the ground that his DNA evidence was found on victim's neck?

The above examples pose challenges to concluding that the evidence automatically establishes that *Defendant* committed the vicious attacks on victim. But to think, on these grounds, that there is no qualitative difference between *DNA* and *Brawl* misses the point. It is really irrelevant whether or not we ultimately think that *Defendant* in *DNA* should be found guilty. The point that I am ultimately making in *DNA* and *Brawl* is that most people, including legal scholars and judges treat cases involving DNA statistical evidence very different from cases involving other types of statistical evidence even if the treatment in DNA cases does not always result in automatically finding for the plaintiff. Sometimes the difference in treatment is seen in their willingness to consider whether the plaintiff in each case has standing. The point is that if we think that *DNA* should be treated more favorably than *Brawl* on the basis of the evidence that is appealed to, then this signals that we view DNA evidence to be qualitatively different from other forms of DNA evidence.

treatment of DNA evidence in contrast to other forms of statistical evidence, the theoretical framework must explain why courts would find *Defendant* guilty in ***DNA*** but not in ***Brawl***.Thus:

DI The theoretical framework must explain why it is acceptable to find for the plaintiff in *DNA* but not in *Brawl*. (To save space, if the theoretical framework fails to satisfy *DNA*, then that will suffice for showing that it does not respect **DI**.)

2.3 ADAPTIVENESS REQUIREMENT

A theory of individualized evidence must not only show what makes DNA evidence unique, it must also show how individualized evidence adapts to different standards that courts might apply. Jeremy Bentham (1843) was among the first to recognize that legal proof "admits of, and exists in, different degrees of strength," where "the practice of wagering affords at the same time a proof of the existence, and a mode of expression or measurement for these quantities or degrees".[39] Undoubtedly, the law imposes different standards of proof, corresponding to the various legal proceedings. The various standards have implications on the role of evidence. Indeed, the same collection of evidence that does not meet the requirements for a guilty verdict in one proceeding may meet the requirements for an imposition of liability in another.

2.3.1 The O.J. Simpson Trials

One historical example of this point can be seen in the civil and criminal trials

[39](See 223)

of Orenthal James Simpson (hereafter, O.J.). After Nicole Brown Simpson's and Ronald Goldman's bodies were found on the steps of a walkway in front of the former's condominium, O.J. was charged with two counts of first-degree murder. In October of 1995, O.J. was found not guilty of all the charges in the criminal case. However, less than sixteen months later, using virtually the same evidence offered in the criminal case, the jury in the civil case found O.J. personally liable for the wrongful deaths of the same victims. What accounts for the different verdicts? The facts were the same; the defendant was the same; even the victims were the same. One explanation for the different decisions in the two cases is that *two different* standards of proof were operative. Specifically, the standard of proof in criminal trials is more stringent than it is in civil trials. That is the reason why, holding fixed stipulated facts of the incident in question, O.J. was found liable of wrongful death in one proceeding but not guilty of double homicide in another proceeding.

2.3.2 A matter of different standards

One way of understanding the various standards of legal proof is in terms of probability. According to this view, burdens such as proof by preponderance of the evidence, proof by clear and convincing evidence,[40] and proof beyond a reasonable doubt[41] correspond to probability thresholds, which must be satisfied to return a

[40] Burdens such as preponderance of the evidence and clear and convincing evidence are used in civil proceedings.

[41] Proof beyond a reasonable doubt is used in criminal proceedings

verdict of guilt or liability.[42] For example, the standard of proof that is supposed to be applied in civil trials is the *preponderance of evidence* standard - also known as the 'balance of probabilities'. To find a defendant liable by this standard, the fact-finder must judge it to be more than .5 likely on the evidence that the defendant is liable. A slightly more stringent proof is the *clear and convincing* standard. According to a study by McCauliff (1982), many federal judges and United States Supreme Court justices believe that to find to find a defendant liable by this standard, the fact-finder must judge it to be approximately .75 likely on the evidence that the defendant is liable.[43] Finally, the strictest standard of proof is *proof beyond reasonable doubt*, which is used in criminal trials. To find a defendant guilty by this standard, the fact-finder must judge it to be at least .95 likely on the evidence that the defendant is guilty.

Obviously, the fact that different legal proceedings employ different standards of proof has bearing on frameworks attempting to explain individualized evidence. It will not do if frameworks can only accout for one standard of proof, but not all of them. This means, then, that a robust framework must show that what counts as individualized evidence adapts to the different standards of proof across the different standards in the law. Put somewhat differently:

[42] (See Smartt; 6, 20)

[43] An informal study conducted by Judge Weinstein revealed that judges in the Eastern District of New York associated this standard with probability thresholds ranging from .6 to .75 for example, and a subsequent study by McCauliff 1982 found that hundreds of judges and United States Supreme Court justices associated this standard with an average probability threshold of .75

DII The theoretical framework must adapt to the different standards of proof that the courts may apply.

2.4 WRONGFUL CONVICTIONS

2.4.1 The initial basis for the third desideratum

The third desideratum deals with inevitable miscarriages of justice. One noticeable feature of each standard reveals that none of them requires that the trier of fact have absolute certainty that the plaintiff has made her case. Indeed, even proof "beyond a reasonable doubt" does not mean that the state proves its case to an absolute certainty. Were such a standard required, then very few criminals would ever be brought to justice.

There is, however, one downside to not requiring absolute certainty - it leaves open the possibility of miscarriages of justice, where a person is convicted and punished for a crime that he or she did not commit. Indeed, in any criminal legal system, a subset of innocent defendants is virtually guaranteed to be convicted. Several philosophers of legal reasoning agree that an innocent defendant can be proven guilty. As Gardiner 2019 writes:

[I]t is a desideratum of an account of legal evidence that - if evidence is compelling but misleading - the burden can be satisfied even if the judgment is false ... an innocent defendant can appear beyond guilty beyond reasonable doubt.

Similarly, Blome-Tillman 2015 argues:

More often than desirable, courts find defendants liable that are in fact not at fault. In some such cases of wrongful conviction, the defendant is found liable on the

basis of very strong but ultimately misleading evidence. In such cases, the court is not at fault. The standards of proof have been met by strong, if misleading evidence. Call such cases *no-fault wrongful convictions.*[44]

2.4.2 Epistemic justification and the burden of proof

In the philosophy of law literature, some have drawn a link between individualized evidence and traditional epistemology. Discussions of the connection tend to take many forms. The most common is to suggest that proposed solutions used in traditional epistemology, particularly, accounts of propositional knowledge, can aide and inform solutions to issues in philosophy of law. For example, as will be discussed in Section 3, Thomson's causal connection and guarantee argument borrows the idea of a causal connection, which is central to Alvin Goldman's causal theory of knowing. Similarly, Enoch and his colleagues' sensitivity argument is based on Robert Nozick's truth tracking theory.

An often overlooked, but no-less significant, connection is that between epistemic justification and the legal burden of proof. The similarities are striking. For instance, epistemic justification is said to be the right standing of a person's beliefs with respect to knowledge. Indeed, on many accounts, it is a necessary condition for (or, constituent of) knowledge. Also, epistemic justification comes in degrees - an agent can have more or less of it. Finally, justification is also a normative concept. Determining whether a belief is justified can indicate whether one *should*, *should not*, or *may* believe a proposition.

Analogously, satisfying the legal burden of proof is also a constituent of a stan-

[44](108)

dard; namely, returning a finding for the plaintiff. Also, similar to justification in epistemology, the burden of proof can come in degrees. It can either be met or not met. Most of all, assessing whether the burden of proof for a proceeding has been met can indicate whether one *should*, *should not*, or *may* find for the plaintiff. Thus, in many respects, satisfying the legal burden of proof in philosophy of law is the analogue to being justified in holding a belief in epistemology.

If the analogy between satisfying the legal burden of proof and being justified in believing a proposition holds, then this has a significant implication regarding the third desideratum: Most epistemologists maintain that an agent can be justified in believing something that is false. That is, one can have enough good reasons for believing some proposition even if the proposition turns out to be false. Simply put, epistemic justification does not entail having a true belief. But if being justified in believing p does not entail the truth of p, then what could be the legal analogue to this? I believe the answer has to do with finding for the plaintiff. Specifically, the legal analogue to being justified in believing something false is finding a defendant liable for good reasons, though he is ultimately not at fault. This then serves as another reason why it should be possible to satisfy the burden of proof even if the judgment is false.

Thus:

DIII The theoretical framework must show that some false convictions ought to satisfy the burden.

(Note: To save space in Section 3, I will only discuss desiderata that the theoretical framework do not satisfy. Thus, if a theoretical framework appears to satisfy the

above desideratum, I will forgo discussing how it meets that condition)

2.5 CODICILS

As my final desideratum, I propose a simplicity constraint, which is only to be invoked if the framework meets **DI - DIII**. This constraint derives its plausibility historically from work by scientists, mathematicians and philosophers who stress that, *ceteris paribus*, simpler theories are better.[45] Hence, simplicity is seen as a theoretical virtue.[46]My condition puts a constraint on how a test may handle po-

[45] Aquinas adopted simplicity as a theoretical virtue when he writes: "We may assume the superiority ceteris paribus of the demonstration which derives from fewer postulates or hypotheses (1945, 129). Kant, following suit, argues that "rudiments or principles must not be unnecessarily multiplied" (1950, 538-9). This is not far removed from Newton, who includes a principle of parsimony as one of his three "Rules of Reasoning in Philosophy" at the beginning of Book III of Principia Mathematica (1972, 398). Also, Lavoisier writes in support of simplicity: "If all of chemistry can be explained in a satisfactory manner without the help of phlogiston, that is enough to render it infinitely likely that the principle does not exist, that it is a hypothetical substance, a gratuitous supposition. It is, after all, a principle of logic not to multiply entities unnecessarily (1862, 623-4). Finally, Lewis writes: I subscribe to the general view that qualitative parsimony is good in a philosophical or empirical hypothesis (See: 1973, 87)

[46] Not all philosophers share this view. Margaret Urban Walker (2007), for example, argues that simplicity constraints are part of a bias that philosophers should avoid. The bias is a part of, what she calls, the theoretical-juridical model of morality (or TJM), according to which moral theories attempt to "represent" the ideal capacity of the well-equipped moral agent, or try to justify its issue, in a codifiable, compact, consistent procedure(s) for generating or justifying action-guiding judgments

(Walker; 2007, 43). By "compact", I take Walker to mean "simple" or something synonymous to it.

Walker's strongest objection against TJM is perhaps the claim that TJM represents aspects of a social order that fosters, among other things, oppression towards women. I do not challenge whether Walker is right about the causal connection between TJM and gender inequality. For all I know, Walker's point may very well be true. Walker's worry with TJM does, however, raise an important question. By linking gender oppression with TJM, she also intends to undermine standard practices associated with TJM, including codifiability, compactness, and consistency. But why should we think that her criticism, if accurate, shows that we ought to reject every aspect of TJM, including compactness as a procedural approach? We do not say, for example, that P is false, or objectionable merely because of its relation with Q, which we think is false or objectionable. Indeed, a lot more would need to be said about why P's being associated with Q gives us reason to reject P. Put somewhat differently, for Walker's argument against simplicity to go through, she would need to show that simplicity is what makes TJM foster gender inequality (or strongly contributes to it). Or, she would need to raise an entirely new difficulty with simplicity that is separate from the problems she has with TJM.

It is unclear whether she has done either adequately. It is also worth mentioning that Walker does not provide much better in terms of an alternative to TJM. Her own positive theory - the expressive-collaborative model, which underscores (i) giving responsibilities to others; and (ii), methodological suggestions, including transparency testing and strong objectivity - collapses the distinction between 'morality as a social phenomenon' and 'morality as what we have reason to do'. On Walker's

tential objections. It is common for philosophers to invent specially made entities and or conditions (hereafter, codicils) designed to block potential objections to their theory.[47]Such codicils, though not necessarily a sign of error, tend to have several undesirable properties.

To begin, codicils often have little connection to the theories they are de- signed to strenghthen. This is to be expected, given that philosophers only invoke such entities to circumvent problems their theories may lead to. The entities would have no place in the theory were it not for the potential problems the entities are designed to block.

Also, codicils do not always reflect reality. One example of this can be seen in view, morality is merely a social phenomenon, which emerges out of negotiation and takes different forms in different societies. But does not regarding morality in this way inevitably lead to moral relativism? This spells trouble for Walker, given that she wants to avoid endorsing relativism. In conclusion, I find neither Walker's comments against simplicity nor her positive theory as reason to reject simplicity as a desirable property for a theory to possess.

[47] The style of reasoning that I have in mind is sometimes known as ad hoc, which is nothing new in philosophy. There is, however, much debate about what is meant by ad hoc. Irving Copi propounds three different senses of ad hocness: (i) an ad hoc hypothesis is one which is specially made up to account for some fact after that fact had been established; (ii) ad hoc is also used to characterize a hypothesis which accounts only for the particular fact or facts it was invented to explain and has no other explanatory power, that is, no other testable consequences; and (iii), ad hoc can be used to denote a mere descriptive generalization. Such a descriptive hypothesis will assert only that all facts of a particular sort occur in just some particular kinds of circumstances, and will have no power or theoretical scope. (See; 1972, 452-453) My second desideratum targets the first two senses of ad hoc reasoning.

Ptolemy's treatment of Mars' retrograde motion in his treatise the *Almagest*.[48] In the 2nd-century AD, Claudius Ptolemy formalized an earth-centered view of the universe - a view scientists since then have rejected. Although Ptolemy's model had predictive power,[49] it also employed epicycles - a geometric model used to explain the variations in speed and direction of the apparent motion of the Moon, Sun, and planets. Ptolemy's usage of epicycles was not original. By the the end of the 3rd century BC, the Greek astronomer Appollonius of Perga had already developed a view of the universe using epicycles.[50] But Ptolemy's use of epicycles to account for Mars' retrograde motion is one of the grounds of an instrumentalist reading given to the Ptolemaic system.[51]

A third concern with using codicils is pragmatic. The more such conditions are postulated the more unwieldly the theories using them may become. The shift from a Ptolemaic model of the universe to Copernicus' sun-centered model in the 16th century also exemplifies this point. Ptolemy's system, with equants and epicycles, was highly complicated, leading many scientists to view it as messy. Epicycles also appear on Copernicus' system, but required many *fewer* epicycles than Ptolemy.[52]

[48] See DWitt; 2004, 117-124

[49] (See Seeds; 1998, 44-49)

[50] There is evidence of epicycles being developed by Hipparchus of Rhodes, who also predates Ptolemy.

[51] (See MacIntyre; 2006, 10; DeWitt; 2004, 75; and Stathis; 1999, 29)

[52] (See The Copernican Model: A Sun-Centered Solar System)

The logic of Copernicus' model made this possible. Placing the Sun at the center with the planets orbiting it easily explains why fewer epicycles would be needed. It also explains why many scientists regarded the Copernican model as elegant.[53]

Given the difficulties with codiicls, we have *pro tanto* reasons to prefer one theory over its rival, if the former has fewer of these kinds of codicils. That is how I shall understand simplicity, and is the central idea in **DIV**.

DIV relates to **DI**, **DII** and **DIII** in the following way: It must not be the case that in order for the framework to satisfy any one of the previous desiderata that it employs codicils whose sole purpose in the test is to satisfy either **DI**, **DII** or **DIII**. Informally, codicils addressing **DI**, **DII** or **DIII** must have independent justification. A codicil having an independent justification can be satisfied by the following test: had either **DI**, **DII** or **DIII** not been a desideratum, then the codicil would still reasonably be a part of the universality test in consideration.

Thus:

DIII. **The test must not invoke specially made conditions designed only to meet DI, DII or DIII**. (Note: Given that the fourth desideratum puts a constraint on how a test may address the other desiderata, I will not discuss whether a test satisfies **DIV**, if it fails to meets all of the other desiderata.)

[53] (See Seeds et *al.*; 2007, 60)

CHAPTER 3

Current Theories

3.1 MORAL THEORIES

An indication that a problem proves difficult to resolve can often be seen by the number of proposed solutions to it. The literature on statistical evidence is extensive, with many attempts to explain the difference between statistical and individualized evidence. Some attempts explain the social and moral issues that would arise if courts began to convict individuals on the basis of statistical evidence alone. Other attempts appeal to epistemic norms as marking the difference between both types of evidence, the idea being that individualized evidence possesses epistemic properties but bare statistical evidence does not.[54] The first collection of arguments seem to view the problem of statistical evidence as a moral issue, and consequently provide a moral theory of individualized evidence. The second collection of arguments, on the other hand, seem to view the problem of statistical evidence as an epistemic

[54] I do not claim that these are the only two categories into which extant theories of individualized evidence fall. There is work on naked statistical evidence that relies on probability. Contributors include Edward K. Cheng (2013), Reconceptualizing the Burden of Proof, 122 Yale Law Journal, and Marcello Di Bello (2019), Trial by Statistics: Is a High Probability of Guilt Enough to Convict? 128 Mind. There is also work on statistical evidence that relies on the idea of "weight". One contributor is Alex Stein's Foundations of Evidence Law, Oxford University Press, 2005. Each category is worthy of attention and consideration. Given space and time constraints, I thought it worthwhile to focus on the two often discussed categories of individuualized evidence.

issue, and thus offer an epistemological framework of individualized evidence. In this Subsection, I evaluate some of the most influential moral proposals in light of the aforementioned desiderata. Although the authors did not design their accounts with the above conditions in mind, it is nonetheless worthwhile to see how their accounts fare with the conditions. The first batch of proposals call attention to the issues and violations that statistical evidence cases cause on the defendant. As such, I will refer to them as defendant-based concerns.[55]

3.1.1 Defendant-based frameworks

3.1.1.1 Defendant-Specific Claims

Mark Colyvan et al., claim that the problem with relying solely on statistical

[55] Contributions in this camp include Posner, Richard (1999), who claims that a resort to statistical evidence reveals that no other evidence could be found, and that this in itself indicates the weakness of the plaintiff's case; Wright, Richard (1988) who argues that eyewitness testimony is about the specific case but merely statistical evidence is not and thus is no way related to determining what happened in the specific case; Colyvan, Mark et al., (2001) who maintain to find for the plaintiff in a naked statistical evidence case is to punish the defendant for being a member of a reference class; Lilquist, Erik (2002) who asserts that to punish the defendant in a naked statistical evidence is to compromise justice in order to achieve a more ecient result in the overall class of cases; Nesson, Charles (1985) who states that rulings based on statistical evidence may be illegitimate and unacceptable, despite any potential probative value of the evidence, because of their adverse effect on public trust in the adjudication system; and Wasserman, David T. (1991) who claims that relying on statistical evidence violates the relevant party's autonomy and individuality, and perhaps even her free will and agency.

evidence concerns what it does to the defendant. Spefically, the defendant is being punished for being a member of a reference class. The authors refer to this problem as the reference-class problem. To be sure, the problem with Blue Bus case and any other naked statistical evidence case isn't so much that the probability of guilt (or liability) is not sufficiently high. Rather, the problem is that that reference-class evidence is not specific to the individual in question.[56]

There are many features to like about the concerns that Colyvan et al. raise as it relates to naked statistical evidence. Indeed, it does appear that to pro-plaintiff findings in such cases punishes the defendant for belonging to the a particular referent class. Further this does seem to be (at least) one of the reasons we are uncomfortable with how such cases are handled. Notice, however, that Colyvan and his colleagues offer no explanation as to how DNA evidence, despite also being a form of statistical evidence, is not vulnerable to the same allegation. In fact, their diagnosis of the problem offers little in terms of how to address the other desiderata. At best, the reference-class problem only provides a diagnosis of the issue.

3.1.1.2 Autonomy Considerations

Another argument against naked statistical evidence cases that looks at it from the vantage point of the defendant is offered by David Wasserman.[57] Wasserman claims that a fair trial should always respect the defendants' autonomy, free will, and other values. Indeed, the defendant should be tried while knowing that her

[56] (See Colyvan et al., 2001)

[57] (See 1991)

individuality remains intact. However, relying solely on bare statistics suggests an unfair prediction of the defendant. It assumes that the defendant was bound to commit the wrongful action they are being accused of, even if they did not *actually* commit it. But such a prediction does not respect the defendant's autonomy.

I do not doubt that defendant's free will and other core values are compromised by relying only on statistical evidence cases. However, this concern does not begin to address DI**, DII, or DIII**? **How are we to understand why** DNA evidence is treated differently from other forms of statistical evidence on Wasserman's account? How does Wasserman's explanation adapt to the different standards of proof in the law? Wasserman account does not respect any of the desiderata.

3.1.1.3 About-Relation Claims

Richard Wright propounds an agument dealing with the harms of statistical evidence from the defendant's point of view. Wright distinguishes between evidence that is genuinely about the particular defendant and evidence that is unrelated to the defendant's matter. On his view, naked statistics are merely accidental groupings. As such, they do not count at all "as proof of what actually happened on a particular occasion."[58]

Wright's concern is well noted. However, this does not explain why DNA evidence is treated differently. In what way is DNA evidence about the defendant but other forms of statistical evidence isn't? Wright's explanation is unable to capture our intuition concerning DNA evidence. Indeed, his issue with naked statistics do not satisfy the other desiderata either.

[58](See Wright; 2001, 1056)

3.1.2 Nondefendant-based frameworks

So far we have evaluated several defendant-based, moral frameworks. Unfortunately, when they are considered in light of the desirable properties stated in Chapter 2, none of them yield good results. I now move on to a new set of frameworks. The new theories explain the issues relying on statistical evidence has on many parties apart from the defendant.

3.1.2.1 Social-Acceptability Claims

Charles Nesson claims that verdicts based on statistical evidence are socially unacceptable.[59] By relying on naked statistics, the legitimacy of the court is at stake. Specifically, the public perception of the court will be tarnished if the court were to countenance naked statistics Statistical evidence shifts the message conveyed by the court from one of certainty to one of risk assessment.[60] In doing so, it expressly states the risk of error underlying the judicial verdict and may thereby weaken the stystem's legitimacy in the eyes of the public.

Unfortunately, Nesson's proposal suffers from the same problem as the frameworks in the previous Subsection. Nothing in Nesson's explanation can be used to provide a guide for what makes DNA evidence different from other forms of naked statistics. Additionally, Nesson's proposal does not seem to handle the other desiderata in chapter 2.

[59] (See Nesson; 1979)

[60] (See, Enoch et al., ,70)

3.1.2.2 Exogenous-Factor Claims

A final argument in the literature is provided by Richard Posner, who claims that resorting to naked statistics reveals something telling. If the claimant has to appeal to bare statistics, then this underscores the tenuity of her case. If this is, indeed, the case, statistical evidence should be accorded less weight simply because it tends to be submitted in circumstances in which the case of the party presenting the evidence is weaker.

I do not doubt that defendant's free will and other core values are compromised by relying only on statistical evidence cases. However, this concern does not begin to explain what makes DNA evidence different. Indeed, it does not satisfy the other two desiderata either.

So far, I have discussed several theories of naked statistical evidence focusing on moral issues. The strength of such theories lies in their pointing out the individual and societal issues with relying solely on statistical evidence. Many of the problems discussed capture our intuitions about the wrong-making features of bare statistics. Despite the virtues of these proposals, it must also be noted that they do not meet the desiderata offered in Chapter 2. I will now consider, what I call, epistemological frameworks. Such theories are thus named, given that they adopt ideas and concepts from theories in the literature of traditional epistemology. When surveying the epistemological theories in the literature, a pattern emerges, not too different from the pattern focusing on the moral issues with naked statistics: the theories within this category can be sorted into one of two subcategories. Some accounts appeal to knowledge, *inter alia*, as the defining feature that separates naked

statistical evidence cases from nonstatistical evidence cases. Such theories will be called knowledge-based accounts. Other approaches appeal to some concept other than knowledge as the defining feature. For convenience, I will refer to this latter category as nonknowledge-based accounts.

3.2 EPISTEMOLOGICAL THEORIES

3.2.1 Knowledge-based theories

3.2.1.1 Causal-Connection and Guarantee Argument

One well-known theoretical framework, defended by Judith Jarvis Thomson (1986), holds that an appropriate causal connection between the litigated facts and the defendant marks the difference between statistical evidence and individualized evidence. Thomson provides three paradigmatic cases, each exemplifiying different forms that the causal relation to the evidence can take. One pathway is, what Thomson calls, forward-looking evidence; another is backward-looking evidence, and a third pathway is one in which the evidence and crime might have a common cause. Table 3.1 encapsulates the pathways

Pathways	**Description**
forward-looking evidence	The evidence caused the crime. *Example*: Not having insurance while driving motivates the perpetrator to leave the scene of a vehicular accident.
backward-looking evidence	The crime causes the evidence. *Example*: Rapidly driving off the scene of an accident causes the perpetrator to leave skidmarks at the crime scene.
common cause	The evidence and crime have a common cause. *Example*: The perpetrator's fleeing the scene of a crime and leaving tire skidmarks indicates a common cause: the perpetrator was driving on a suspended license.

Table 3.3.1

What makes having an appropriate causal connection to the evidence so valuable?

It is that such a connection provides a guarantee that the claim (or charge) is true. For example, the testimony of an eyewitness can prove that a defendant shoplifted, since the fact that the defendant shoplifted is a cause of her testimony. By contrast, the fact that most teenagers shoplift during the week between the hours of 4 and 5 pm is causally unconnected with the fact that any particular defendant shoplifted, and so the former fact is insufficient to prove the latter, even by a preponderance of the evidence.

Thomson's account, however, runs afoul when dealing with **DII** and **DIII**. Beginning with **DII**, Thomson's notion of guarantee seems overly stringent for understanding standards lower than beyond reasonable doubt. As Redmayne 2008 notes, the civil standard of proof cannot require knowledge, since the subjective component of this standard is merely that the trier of fact has greater than .5 credence that the defendant is liable. By contrast, knowing that a defendant is liable requires the trier of fact to fully believe that the defendant is liable, which is a much stronger constraint.

Acknowledging the above problem that the causal connection as guarantee view faces, Thomson extends her account from criminal to civil burdens. According to Thomson, different standards of proof correspond to how sure fact-finderss are of having a guarantee of culpability. She writes,

Our law requires the jury in a criminal case to be sure beyond a reasonable doubt that the defendant is guilty before imposing liability on him; the friend of individualized evidence may be taken to say that the jury must be sure beyond a reasonable doubt that the defendant is guilty because of being sure beyond a reasonable doubt that there are facts available to it which guarantee that the defendant is guilty. Our

law requires the jury in a case in tort to believe no more that that it is more probable than not that the defendant is guilty; the friend of individualized evidence may be taken to say that the jury must believe it is more probable than not that the defendant is guilty because of believing it more probable than not that there are facts available to it which guarantee that the defendant is guilty.[61]

A charitable way to interpret Thomson's treatment of the "clear and compelling evidence" standard is as follows: to meet the standard, the trier of fact must believe she has clear and compelling evidence that there are facts available to her which guarantee the defendant is liable. To satisfy a "reasonable suspicion" standard the trier must believe she has reasonable suspicion that there are facts available to her which guarantee that crime is underway.

I maintain that there are still problems that Thomson's revised version engenders. First, what does it mean to possess clear and convincing forms of evidence that guarantees liability? This is woefully unclear. Secondly, Thomson's revised position still seems overly stringent. However, we understand what it means to possess less than certain forms of evidence that guarantees liability, it seems possible for the fact finder to possess clear and convincing evidence that indicates liability short or guranateeing it.

Finally, Thomson's account does not satisfy **DIII** either. Her knowledge-based account delivers factive standards of legal proof. Thus, on her view, an innocent defendant cannot be proven guilty. But, as was argued in Section 2, this result is unacceptable.

[61](See Thomson; 1986, 215)

3.2.1.2 Probabilistic-knowledge Argument

The next framework, defended by Sarah Moss, is specifically designed to handle the adaptive requirement. Similar to Thomson's causal connection, Moss' view holds that legal proof requires knowledge. But, unlike Thomson, the kind of knowledge that Moss builds her legal theory on is probabilistic, not propositional.[62] To understand probabilistic knowledge, we must get clear on what it means, on Moss' account, to know a probabilistic content.

Moss takes probability spaces to be the content of probabilistic knowledge, where a probabilitiy space is a tuple $s = \langle \Omega_s, F_s, \mu_s \rangle$, such that:

- Ω_s is a set of possible worlds
- F_s is a σ-algebra such that for all $\omega \in \Omega_s$, $\{\omega\} \in F_s$
- μ_s is a probability measure on F_s

Informally, a probability space is a probability distribution over a given set of possible worlds. To see this, suppose the gambling statuses of Sally and Amy are the only facts that interest us. Then there are four possible worlds: one in which both gamble, one in which neither gamble, and two more in which only Sally or Amy gambles. Suppose that I am completely sure that Sally does not gamble, but completely unsure whether Amy does. On Moss' view, my epistemic state is represented by a set of probability spaces, with one probability space for each probability distribution over the two worlds at which Sally does not gamble.

[62] Moss' probabilistic knowledge is not merely propositional knowledge about probabilities. She also denies that it is simply evidential or epistemic probability.

Moss' notion of probabilistic knowledge has several advantages. For example, we can get probabilistic knowledge[63] in all the familiar ways of getting propositional knowledge - *viz*, by testimony, perception, inference, memory, and *a priori* reflection.[64] Second, probabilistic knowledge seems compatible with a recent movement in epistemology: knowledge first, which originates with Williamson (2000). As Patricia Rich writes, "The starting point of knowledge-first epistemology is that knowledge is unanalyzable, more fundamental than belief, and more important than belief."

A third advantage to Moss' probabilistic knowledge has direct application to **DII** - *viz*., it can capture what it means to satisfy the burden of proof across the various legal proceedings. Specifically, on her account, a defendant is proved liable *by preponderance of the evidence* only if the fact-finder has greater than 0.5 credence that the defendant is liable, and that credence constitutes knowledge. Similarly, a defendant

[63] Moss also tells us what theory is not saying. Probabilistic knowledge is not just knowledge of objective chances - *viz*, the objective chance that someone is a gambler is either 0 or 1, but my credences are intermediate. Similarly, probabilistic knowledge is not propositional knowledge about probabilities. Also, we do not have probabilistic knowledge in a trivial sense, just because we can derive probabilistic statements from propositional ones. For example, If I can know that Sally probably gambles, the explanation is not (necessarily) that I know that Sally gambles, so I know she probably does. One can know that something is the case without grounding her knowledge in the fact that it is the case. Finally, probabilistic knowledge is not a second kind of knowledge. Moss holds that all knowledge is fundamentally probabilistic and familiar propositional knowledge is just the limiting case in which the proposition receives probability 1. The limiting case is, according to Moss, "nominally probabilistic knowledge. (See, Moss; 2017, Ch. 1.4) This is how probabilistic knowledge can therefore be explained in just the same way that we would explain traditional propositional knowledge.

[64] (See Moss; 2018, 86)

is proved liable *by clear and convincing evidence* only if the judge or jury knows an even stronger probabilistic content, and guilty *beyond a reasonable doubt* only if the judge or jury knows a still stronger probabilistic content. This view generalizes to saying that a standard of proof is met when the fact-finder knows a specific content - namely, the set of probability spaces according to which the probability of guilt or liability meets or exceeds the relevant threshold for the standard.

Finally, Moss' probabilistic knowledge account of legal proof can show why mere statistical evidence is usually insufficient for probabilistic knowledge. Recall ***Prison Riot***, which is stated in Chapter One. According to probabilistic knowledge, the mere statistical evidence in this case fails to provide the fact-finder with probabilistic knowledge. Although it might justify a 0.96 credence in the claim that any individual prisoner committed the murder, the fact-finder cannot know that Smith is probably guilty, since the mere statistical evidence does not allow her to rule out a relevant alternative; namely, that Smith is the one exception within the reference class. Moss argues that not only does this solution show why the evidence in ***Prison Riot*** fails to provide probabilistic knowledge that Smith is probably guilty, but it can be extended to an explanation of the insufficiency of mere statistical evidence in other proof contexts. Although the strength of the probabilistic content will vary as the standard of proof varies, the basic conditions on probabilistic knowledge remain fixed. So even in proof contexts involving much lower probabilistic thresholds, the mere statistical evidence might provide justification for a certain credence but not knowledge of this content.

Turning now to the above desiderata, I do not think that Moss' probabilistic knowledge satisfies **DI** or **DIII**. Starting with **DI**, if ruling out all relevant alterna-

tives is required for knowledge, then it is unclear how probabilistic knowledge can account for the exceptional treatment given in DNA cold-hit cases. While a very large number of relevant alternatives is eliminated in such cases, it is also true that not *all* relevant alternatives are eliminated. Thus, *prima facie*, Moss' proposal does not adequately reflect the deliberation that the fact-finder *actually* employs in a sizable portion of legal cases where statistical evidence suffices for a verdict.

It is worth pointing out that textual evidence seems to suggest that Moss is mindful that her account cannot explain the exceptional treatment of DNA evidence. In setting out the scope of her project, she writes: "To clarify what I am aiming to explain: our hostility to statistical evidence is not without qualification. Courts have also been increasingly willing to issue conviction on the basis of random match DNA evidence."[65] This then leads her to conclude that "the problem of statistical evidence is not the problem of explaining why statistical evidence is never sufficient for a verdict, but rather why it is insufficient in many normal cases." A natural reading of Moss' words suggests that she is claiming that her account of legal proof is aimed at showing why statistical evidence is insufficient in "many normal cases." This would seem to imply that DNA statistical evidence cases are, at best, *abnormal*. More troubling, Moss' words also suggest that her account is not designed to explain why courts have been increasingly willing to issue conviction on the basis of random match DNA evidence.

If my analysis of Moss is correct, I think that there are two problems that her claims invite. To begin, we must ask whether it is still appropriate to suggest that the law's favorable treatment of DNA cold-hit cases is abnormal? While it may have

[65](See, 2021, 13)

been appropriate to suggest this in the early 80s when many courts began to accept DNA evidence, this classification is no longer apt. As Moss, herself, notes, "Courts have been increasingly willing to issue convictions on the basis of random match DNA evidence."[66] The current direction the law seems to be going suggests that what was once abnormal is now the norm.[67]

A second, and perhaps deeper, worry relates to **DI.** If a robust theory of individualizing evidence ought to explain what makes DNA evidence different from other forms of statistical evidence, then simply suggesting the former as "abnormal" would not constitute an explanation. Indeed, the question is just pushed back one step further: We began asking "*What makes DNA statistical evidence cases exceptional?*" But now the question is *"What makes DNA statistical evidence cases abnormal?*" Nothing is gained by designating DNA evidence as abnormal.

Given my foregoing concerns, we might be tempted to think that a plausible rejoinder is available to Moss. Instead of concluding that she thinks of DNA cases as "abnormal", she could just be setting aside cases involving DNA evidence and asking why statistical evidence is not sufficient for conviction in normal cases. In fact, if pressed, Moss could say that what makes DNA cases unique is that DNA evidence alone generates the kind of probabilistic knowledge that satisfies the relevant standard of proof; whereas, other naked statistical evidence alone won't generate probabilistic knowledge. But this line of argumentation also invites further diffi-

[66] (See 2021, 14)

[67] See Roth 2010 for discussion of the recent rise in "pure cold hit" prosecutions.

culties. How is it that DNA statistical evidence cases alone generate probabilistic knowledge? What properties do such cases possess that other statistical evidence cases lack? Moss would owe us an explanation of this. Moreover, by looking at her explanations of why other forms of naked statistical evidence cases fail to generate probabilistic knowledge, I do not see how her explanation can be anything but *ad hoc*.

Probabilistic knowledge does not seem to respect **DIII** either. Or, it is not able to solve it without great difficulty. To see this, note that given the fact-finder is required to probabilistically know that the defendant is liable (or guilty) and since factivity[68] is a constituent of knowledge, requiring that the fact-finder knows that the defendant is liable (or guilty) implies, as the objection goes, that it actually be the case that the defendant is liable (or guilty). However, it seems that there is at least one response that Moss can give to this objection. The objection seems to conflate Moss' notion of probabilistic knowledge with propositional knowledge. The latter requires factivity but the former does not - or, at least, it does not seem to require what we traditionally mean by factivity. Thus, the factivity implication leveled at probabilistic knowledge is unwarranted.

I believe that there are at least two problems with this response. The first is epistemic. Given that factivity is the least controversial property of propositional knowledge, and hence not a property that is easily abandoned in moving to a probabilistic concept, probabilistic knowledge faces a severe interpretive challenge. To

[68] Factivity means that an agent can only know some ϕ when ϕ is true, or ... when ϕ. Put simply, knowledge entails truth.

offer a notion of factivity to fit probabilistic knowledge, Moss rephrases it in terms of probability spaces:

> The set of probability spaces according to which Smith knows that Jones probably smokes is a subset of the set of probability spaces according to which Jones probably smokes, and that is why the former entails the latter. The same goes for any subject S and any probabilistic content p. The probabilistic content that S knows p is a subset of the content p itself. At a second pass, that is what it means to say that probabilistic knowledge is a factive attitude[69]

There are, however, severe challenges with this interpretation. Rich writes "On the face of it ... a probabilistic ϕ cannot be non-trivially true. Jones either smokes or she does not; there is no possible world in which she probably smokes." Moss insists nonetheless that probabilistic ϕ can be true, as required for probabilistic knowledge. She relies on our intuition, especially about the assertability of probabilistic statements. While there seems to be some linguistic data on her side, the notion of probabilistic knowledge cannot do much work in epistemology until there is a more concrete method of determining (and modeling) when an agent has probabilistic knowledge of some ϕ, which requires determining ϕ.[70]

[69] (See Moss; 2017, 126)

[70] (See Rich; 2020)

Secondly, there is textual evidence that Moss acknowledges that her view implies factive standards of legal proof. In defending her view against possible objections, she writes:

> "The first objection, which is perhaps the prevailing objection in philosophical discussions of legal proof, is that the knowledge account delivers factive standards of legal proof. According (*sic*) my account, an innocent defendant cannot be proven guilty. Some find this result unacceptable."[71]

If Moss' proposal does not imply factive legal standards of proof, then we would expect her to correct the errant accusation. However, this is in fact what she does not do. What we instead read is a defense of factivity in legal standards of proof. One last piece of textual evidence that Moss' account implies factivity can be seen in a footnote where Moss attempts to deal with the objection. She writes: "For simplicity, I focus on the criminal standard in this section. The civil standard of proof by a preponderance of the evidence is also factive, in the sense that a defendant cannot be proved probably liable unless the defendant is probably liable."[72] If factivity were not a consequence of her proposal, it would seem odd for her to write this footnote.

I conclude then that Moss' framework has many attractive features. For example, the probabilistic account is able to address several legal puzzles.[73] Addition-

[71] (See Moss; 2018, 27-28)

[72] (See Moss, 2018, 27-28)

[73] For example, in liability cases, it explains why we think it is not permissible to

ally, Moss' solution to the problem of mere statistical evidence generalizes beyond legal philosophy. Moss handles statistical generalization in social life in a similar way.[74]However, the underlying issue with Moss' account is her requirement that the fact-finder be in a position of knowing that legal burden is satisfied. Specifically, requiring that the fact-finder eliminate all relevant alternatives does not adequately reflect the fact-finder's deliberation in determining outomes of cases. This is, in part, why her account cannot provide a non *ad hoc* explanation for why courts sometimes do find statistical evidence sufficient for a verdict. To overcome this problem, a robust account of individualized evidence must not only explain why courts do not find statistical evidence sufficient for a verdict in some cases, but also why they do in other cases.

3.2.2 non Knowledge-based theories

Knowledge-based theories possess several strengths. However, a common problem, as it relates to the desiderata proposed in Chapter Two, is that such theories seem ill-suited for satisfying **DIII**. Even with a nontraditional notion of knowledge, as is found

find for the plaintiff in the blue-bus case. Similarly, in criminal cases, it explains why it is impermissible to return a guilty verdict to a random defendant in **Prison Riot**.

[74] Indeed, in cases of social generalization where one's evidence consists only of mere statistical evidence, we are unable to rule out the possibility that an individual person is unlike an arbitrary member of their reference class. Thus we will not be in a position to *know* that the individual probably has the target property.

in Moss' probabilistic knowledge, such views have difficulty satisfying **DIII**. The next set of theories that I will discuss in this chapter will, similar to the preceding two theories, appeal to epistemic properties in marking the distinction between statistical and individualized evidence. But, unlike the above theories, the epistemic property appealed to does not entail or assume some form of knowledge.

3.2.2.1 The Reasonable and the Relevant Argument

The first theory of individualized evidence, under this category, that I will consider is defended by Gardiner (2019), who appeals to a relevant alternative account of knowledge, according to which a subject S knows that p only if she can rule out relevant alternatives to p. By an alternative, Gardiner has in mind an error possibilitiy that is best understood as a collection of sub-alternative possibilities. Such possibilities are incompatible with the target belief. Gardiner writes, "If the alternative obtains, then the belief is false."[75]

Not all possibilities need to be considered, let alone ruled out. Some error possibilities may rightly be ignored, because they are, for example, farfetched. Following Lewis (1996), Gardiner defines relevant alternatives as those error possibilities that cannot be properly ignored.[76] Such alternatives, themselves, will be ones that are important, reasonable, and relevant. Finally, to rule out an alternative is to possess some evidence that discriminates the truth from the alternative.

This relevant alternatives framework for knowledge is then used to explain the

[75] (See Gardiner, 2019, 293)

[76] (See Gardiner, 2019, 295)

error possibilities that can be disregarded in each legal standard of proof. Gardiner maintains that the error possibilties can be classed into three concentric circles; three tiers of disregardability. Starting with the preponderance of evidence standard of proof, this standard encloses all preponderant error possibilities. Such error possibilities are the most important and significant and so the least disregardable.

A second tier - that is, the clear and convincing standard of proof - marks those error possibilities that must be addressed in order to satisfy the clear and convincing evidence standard.[77] One's body of evidence might satisfy the preponderance standard, by addressing the most preponderant error possibilities, but fail to satisfy the more demanding standard. The final standard of proof - the proof beyond reasonable doubt - requires addressing all reasonable error possibilities. Gardiner writes, "This is a more expansive set than the alternatives relevant to clear and convincing evidence."(See Gardiner, 2019, 299).

Gardiner then sets forth a relevant alternatives condition on legal standards of proof in terms of error possibilities that can be disregarded.

- General: Claim p is established to a legal standard L only if the evidence adduced rules out the L-relevant error possibilities.
- Preponderance: Claim p is established to a preponderance of the evidence only if the evidence adduced rules out preponderant error possibilities.
- Intermediate: Claim p is established to the "clear and convincing evidence" standard only if the evidence adduced rules out the error possibilities relevant to the clear and convincing evidence standard.

[77](See Gardiner, 2019, 299)

- Reasonable doubt: Claim p is established beyond reasonable doubt only if the evidence adduced rules out the reasonable error possibilities.

Moving onto the desiderata, I believe that Gardiner's proposal is unable to satisfy **DI**, or it at best is not able to do so without great difficulty. I begin by applying the aforementioned litmus test to Gardiner's reasonable doubt condition. To explain why the fact-finder is to find for the plaintiff in ***DNA*** but not in ***Brawl***, Gardiner would be committed to saying that the evidence adduced in ***DNA*** (and all *DNA*-like cases) rule out all reasonable error possibilities, but the evidene in ***Brawl*** does not. But, this is not very helpful. The issue here seems to be one of vagueness. A deeper look into both cases reveals that there are similarities that do not seem to be explained by Gardiner's account. In ***DNA*** and ***Brawl***, there is a high degree of certitude in the defendant's guilt. Moreover, both cases have individuals who are not ruled out. Given these facts, Gardiner would have us believe that the reason we think the cases should be treated differently is that the evidence adduced in ***DNA*** rules out the reasonable error possibilities, but not in ***Brawl***. But it is difficult to see what reasonable error possibilities could account for the different treatment in each case. At the very least, Gardiner owes us an explanation of how her theory would make sense of this.

3.2.2.2 Sensitivity and Incentives-Based Argument

An older but well-known theory of individualized evidence is defended by Enoch

et al (2012) and Enoch and Fisher (2013)[78], who propose, what may be called, a hybrid approach, which possesses both epistemic and practical features. Beginning with epistemic considerations, Enoch et al (2012) and Enoch and Fisher (2013) base their view on a key component of Robert Nozick's (1981) truth tracking theory of knowledge, which asserts the following:

Sensitivity of Beliefs (SB):

S's belief that p is sensitive just in case if p were not true, then S would not believe that p.

The authors develop a notion of evidence based on SB.

Sensitivity of Evidence (SE):

Evidence e is sensitive to p just in case if p had been false, then e would have been false.[79]

The authors then use SE to explain the distinguishing feature between individualizing and statistical evidence. The reason statistical evidence is unreliable is that the evidence is insensitive - the evidence obtains even if the crime would not have. Individualized evidence, on the other hand, is sensitive - if the defendant did not

[78] The sensitivity solution has stimulated the debate on naked statistical evidence, and been investigated in terms of a Bayesian likelihood ratio by Edward Cheng (2013; 1269-1270) and Marcello DiBello (2018;10)

[79] The idea of epistemic sensitivity is developed to compare possible worlds. But it should be noted that the idea uses the term 'sensitivity' in a way that differs from the statistical concept of sensitivity. In statistics, for example, *sensitivity* refers to the probability of the evidence given that the hypothesis is *true*, but Enoch et al., talk about 'sensitivity' with regard to the possibility of seeing the evidence given that the hypothesis is *false*.

commit the crime, the individualized evidence would not have obtained.

Despite the epistemic value sensitivity derives from distinguishing between individualized and statistical evidence, Enoch et al., argue, however, that sensitivity does not have legal value and consequently should not be of legal concern.[80] To provide a distinction that is of legal concern, the authors get inspiration from an issue that Sanchirico (2001) has raised against character evidence.

Character evidence is used to describe a character trait of a person, such as a tendency to tell lies.[81] It is often thought of as a type of interpersonal statistical evidence, since it does not directly link the relevant individual to the wrongful act of which he or she is accused. In fact, as a general rule, character evidence may not be used in the guilt phase of a trial; it is only admissible in the sentencing phase. Sanchirico argues that the use of such evidence in the sentencing phase provides no incentive for previously convicted persons to refrain from wrongdoing.[82] Evidence law should be, among other things, concerned with supplying good incentives for behavior of agents outside the legal process.

Enoch et al., believe that Sanchirico's concern has application to naked statistical evidence cases: if bare statistics suffices for legal faultfinding, irrespective of whether the accused actually acts unlawfully, then what reason is there for the accused to

[80] (See Enoch et al. 2013: 211-215; Enoch and Fisher 2013: 577-581)

[81] It could also include a witness testifying that a party in the case has a reputation in the community of being argumentative.

[82] (See: 2001, 1227 -29)

refrain from the unlawful act? Obviously, little to none, because "whatever he decides will have negligible influence on the likelihood of his being punished."[83] For convenience, I will refer to the view adopted by Enoch et al., as incentives-based explanation or incentives, for short.

Turning to the **DI - DIII**, does the incentives-based account satisfy all of them? I believe the answer is ultimately *no.* Similar to the causal account, the incentives-based account does not meet the final two desiderata . Starting with the adaptive requirement, the incentives-based explanation seems silent with respect to the shifts in burden of proof. Indeed, it appears more as a one-size-fits-all approach. This belief lends itself to concluding that the incentives-based account does not respect **DII**.

We may, however, think that the problem is easily remedied were Enoch et al., to adopt a proportional incentives-based view, according to which the degree of incentives that is needed to suffice for meeting the burden of proof varies in accordance with the legal case at hand. For instance, in criminal proceedings, the degree of incentives necessary to find for the plaintiff would be much higher than what is needed to find for the plaintiff in civil matters. Going this route, however, is anything but straightforward. Indeed, the authors would need to answer several questions: How would incentives be measured? Whose incentives matter - the accused or the general public? What degree of incentives suffices for satisfying the burden of proof in civil cases? There are also civil cases in which it would seem difficult to even implement incentives as a criterion. For example, how would the incentives-based approach

[83] (Enoch and Fisher 2013; 583)

work in settling right to die hearings?

The problems the incentives-based explanation has with **DII** seem to extend to **DIII**. In claiming that the removal of inecient incentives from lawful conduct is the defining feature of individualizing evidence, Enoch et al., do not quantify how much disincentives for lawful conduct need to be removed to satisfy the burden of proof. Thus, it is at best unclear whether some false convictions satisfy the one-size-fits-all approach.

In sum, the incentives-based explanation cannot adequately explain individualized evidence. While it may satisfy **DII**, it fails to properly address **DI** or **DIII**.

Of special interest, however, may be whether the sensitivity account of individualized evidence, also adopted by Enoch et al., fares any better. While the authors do not think that sensitivity is of legal concern, it still seems worthwhile to explore how it stands with respect to incentives. In the next Subsection, I explore whether Sensitivity - Enoch et al's other proposal - fares any better than the incentives-based account.

Prima facie, the incentives-based explanation, similar to sensitivity, seems to satisfy **DI** - at least, Enoch et al., would have us believe that it does. To see this, consider how it handles ***DNA***, which is stated in Section 2. Assuming no further evidence besides what is uncovered at the crime scene emerges, most people believe that the DNA sample found at the scene is enough to return a guilty verdict. Enoch et al. believe that Sensitivity can show us why that is. On the Sensitivity account, we would have to ask, "Were *Defendant* innocent, would the trier of fact still have convicted him? A natural way to interpret this counterfactual question is as follows: "Had *Defendant* not been guilty but the DNA evidence nonetheless matched his

DNA, would the trier of fact still have convicted him?" The answer seems to be *yes*. However, the authors do not believe that this is the correct way to evaluate the question. Following Lewis' (1986) possible worlds semantics, Enoch et al., argue that the possible world in which *Defendant* is innocent and his DNA sample does not match the DNA at the crime scene (or, sample-clashing world) is *closer* to the actual world than the world in which he is innocent but his DNA sample does match the DNA at the crime scene (or, sample-matching world). The authors, then, propose another way to formulate the relevant question: Had *Defendant* not been guilty and the DNA evidence did not match his DNA, would the trier of fact still have convicted him? The answer here, it seems, is *no*.

Problems, however, emerge for Sensitivity when we give a closer look at Lewis' world similarity metric. Lewis maintains that worlds are similar to the actual world the fewer miracles or violations of the actual laws of nature they possess. Furthermore, no world is closer to the actual world than the actual world itself. This last condition poses a problem for Sensitivity. For there are worlds in which sample-matching obtains: cases in which the evidence collected is contaminated; cases in which evidence is planted.[84] This means that, on Lewis' approach, there will be some sample-matching worlds that are closer to the actual world than sample-clashing worlds. Such sample-matching worlds are ones in which *Defendant* is innocent, de-

[84] Take for example, the case of Adam Scott, whose DNA matched with a sperm sample taken from a rape victim in Manchester, which Scott lived more than 200 miles away and had never visited. Non-DNA evidence subsequently cleared Scott. The mixup was due to contamination - a careless mistake in the lab where the plate used to analyze Scott's DNA from a minor incident was accidentally reused in the rape case.

spite his sample being found in the crime scene. This further means that the authors' proposed relevant question will not work in all *DNA*-like cases. Thus, even if we conclude that Sensitivity handles *DNA* correctly, it seems to do so luckily. At the very least, Sensitivity does not offer much reason to think that it adequately handles cases of DNA evidence in general.

Moving on to **DII**, things do not get any better for Sensitivity. Indeed, it suffers from the same problem as that of the incentives-based explanation: neither admits of degrees. The criterion used to meet the burden of proof in civil proceedings is the same needed to meet the standards in criminal proceedings. But this standard is too burdensome for civil cases.

Finally, Sensitivity does not satisfy **DIII.**. This is because, no false conviction is sensitive. To see why, suppose Nevin is accused of breaking into his ex girlfriend's home in order to retrieve the jewelry he left in her home. The prosecution argues that a neighbor who had met Nevin before saw a person with a similar build and facial features as Nevin break into the victim's home on the date in question and then get into Nevin's car and leave. The prosecution also produces a correspondence between Nevin and the victim where the former vows to take back every item he has given the latter. However, as it turns out, Nevin's twin brother, Neal, actually broke into the apartment. But given their striking resemblance, it was easy to confuse them. Neal also borrowed Nevin's car to commit the crime. We now have a case of a false accusation despite there being compelling, misleading, inculpating, individualized evidence. To determine whether the belief held by the trier of fact is sensitive she must ask, "Were Nevin innocent, would the trier of fact believe Nevin is innocent? Since the defendant is actually innocent, the answer is *no*. Thus, if the burden of

proof requires that the judgment is sensitive, no false conviction satisfies the burden. However, this result is unacceptable.

3.2.2.3 Normic Support Argument

In his work on individual inquirers, Smith (2016) develops a concept of justification featuring a notion he calls *normic support*. Here I provide a simplified introduction to his theory.[85]

According to normic theory of justification, an agent is justified in believing p only if p is normically supported by her evidence."[86] Smith defines normic support as follows: "A body of evidence e is said to normically support a proposition p just in case the circumstance in which e is true and p is false requires more explanation than the circumstance in which e and p are both true." [87] More precisely:

S's belief that some proposition p given her evidence e is normically justified only if her belief that not p given her evidence would call for some *special* explanation - the need for more explanation than the circumstances in which e and p are both true.

So, while it might be very unlikely that my reliable car does not start when I tried to start the engine, the car failing to start does not call for some special explanation. Consequently, my belief that the car does not start would not be normically sup-

[85] For a full discussion, see Smith (2016).

[86] (See 2016; 177)

[87] (See 2016; 40).

ported. On the other hand, my compact car getting low on fuel normically supports going on a cross-country road trip, since some special explanation would be called for, if it never got low during the entire trip.

Several clarificatory points seem necessary. First, Smith states that his normic standard is not meant to be applied to each individual item of evidence presented in a trial, but is to be applied to the total evidence presented.[88] Second, Smith insists that normic support is a *necessary* condition for acceptance of an affirmative verdict.[89] Finally, normic support should not be seen as a rival or alternative to purely probabilistic standards of proof. Rather, normic support should be placed alongside such standards.[90]

The last point has bearing on our present purpose. Indeed, Smith argues that combining his normic standard with current purely probabilistic standards of proof in the law will be compatible with the overall standards of proof being higher in criminal trials. Thus, it seems that Smith's combination is aimed at satisfying the adaptive requirement. Unfortunately, Smith does not say much on how normic support is to be combined with probabilistic standards of proof. The closest he comes to spelling out the adaptive feature of his proposal is by saying that this can be achieved "by conjoining normic support with a more demanding probabilistic standard for criminal

[88] (See Smith; 2018, 1210)

[89] *Ibid*

[90] *Ibid*

trials and a less demanding probabilistic standard for civil trials."[91]

Despite Smith's limited discussion of just how normic support is adaptive, it is still worth exploring whether his proposal respects the desideratum. Relying on what Smith does say, the plausibility of his proposal seems to rest, among other things, on two factors. One factor is the worlds-similarity metric Smith employs to determine what makes one world more normal than another. Smith does address this in Smith (2016).

'Normality', according to Smith, is tied to one's tendency to look for further explanations: abnormal events require an explanation; normal events carry no such demand. For example, it is normal for someone to win the lottery. When we hear that someone won, we do not look for a further explanation. However, matters are different when we hear that someone told a lie. Lies tend to require an explanation - it is a departure from normal circumstances.' (See Smith, 2016)

Smith's normality view seems to have several advantages. It provides a response to the lottery puzzle. [92] It respects multi-premise closure[93] along with a host of other formal principles. However, there are several issues the normality view invites as it relates to **DII**. One worry about the overarching picture is that normalcy is tied too closely to our inclinations to look for explanations. It is not clear that our inclinations are fit to do the job required. Even if the normic conception does well

[91] (See Smith; 2018, 1211)

[92] (See, Smith 2016; chapter 3).

[93] (See, Smith 2016; chapter 4)

in a select range of cases, as McKenna notes, "there will be cases where we have no clear guidance regarding application of the notion. Asking,

Which event is more normal?' or 'Would we look for an explanation? will inevitably not always be enough to go on. Additionally, Smith provides no advice for navigating disagreements about what counts as normal. For example, is winning the lottery *more normal* than lying? Is it more normal for all 100 invited guests to attend the party than for someone to stay home? I do not share all of Smith's intuitions about what is most normal in these cases, and I expect these kinds of difficulties will only be multiplied when making judgments about comparative normality, which is needed to provide a degreed notion of justification. Moreover, our inclinations are subject to biases, as Smith rightly notes. A notion of justification that tracks those inclinations too closely will be prone to inherit those biases.

Another factor is a specification of what it means for one probabilistic standard to be more demanding than another. Unfortunately, Smith does not address this here or elsewhere. At first blush, we might think that for one probabilistic standard to be more demanding than another simply means that the former has a higher threshold for meeting a burden than does the latter. But, even on this conception, there are still some details of the relationship between normic support and varying degrees of probabilistic standards in the law that Smith still needs to work out. Here are just a few questions: What does it mean for normic support to be placed alongside probabilistic standards? Is normic support an additional requirement for satisfying a legal burden of proof? Can one's total body of evidence normically support her legal claim and yet not satisfy the corresponding probabilistic burden of proof? The issues with normic support suggest that we look elsewhere for a framework that satisfies

our desideratum.

3.2.2.4 Reasonable Conviction Argument

Littlejohn (2017) has argued that the difference between the two types of evidence is that a belief should be judged against the standard of acquiring knowledge, which statistical evidence cannot meet. Fundamental to his view is a condition for what is required to punish:

Reasonable Conviction (RC): It is not permissible to punish a defendant if it isn't reasonable to believe the defendant to be guilty.[94]

The claim made is that it is permissible to punish someone only if it is reasonable to believe that he/she is guilty of the charge. Littlejohn writes:

My defense of Reasonable Conviction begins with a reminder that punishment is an act that differs in an important way from acts like betting on football matches or taking umbrellas. This is because punishment is supposed to be a way of holding someone accountable and it involves a backwards-looking element that other actions lack. Thus, the act in question (e.g., imposing a prison sentence) has to be guided by certain kinds of consideration to be a punishment... If we have a system of rules that governs decisions to punish or to refrain from punishing, it would seem that the rules should require that the decision to impose the harms associated with punishment be made only when the punishment can properly express blame or at least treat the defendant as accountable for some specific deed. It would not be proper to blame unless the relevant parties could properly believe that the defendant did something blameworthy.

[94](2017; 15)

Reasonable Conviction (or, RC) explains the implications of holding someone criminally responsible for alleged wrongdoing. If we are willing to exact punishment on the defendant, then we ought to be willing to express blame to them for a specific deed. This seems right. Thus, RC is an appropriate expression of punishment in criminal proceedings. But for RC to satisfy the adaptive requirement, it must also be the case that RC (or something similar) is also what courts use in civil (and other legal) proceedings. Given that RC is primarily concerned with reasonable grounds of punishment in criminal proceedings, this already suggests that it may *not* be adaptive to the other standards across legal proceedings. Yet, it is still worth exploring the merits of this hunch.

Turning to the desiderata, I do not think that (RC) satisfies **DII**. While requiring that it is reasonable to believe that the defendant is guilty in order to convict seems appropriate in criminal proceedings, it does not jibe with a balance of probability standard (or even a clear and convincing standard) of proof for a couple of reasons. One reason rests on the difference between believing that something *is* the case versus believing that the claimant has met a standard of proof . Suppose that while turning at an interection, I am T-boned by another driver. We both blame each other for being at fault. Given our impasse, we decide to resolve the matter in small claims court. In general, the fact-finder has one of two ways of ruling. She may either judge that one of us is liable and thus must pay for damages of both vehicles. Alternatively, she may judge that the evidence does not show who is at fault in which case both complainants would be responsible for repairing the damages done to their respective cars. Suppose I am found at fault. Surely, this does not mean that the fact-finder fully believes, or even thought it reasonable to believe, that I am at fault.

Given that the fact-finder must have greater than .5 credence that I am at fault, it is possible that she believe that the evidence on balance suggests I am at fault and yet suspend belief as to whether I actually am at fault. Indeed, the law does not require that the fact-finder have full, robust conviction in the plaintiff's account of events whenever finding the defendant liable. One reason for this is that the law imposes rules to encourage that the fact-finder operates in a thoroughly binary fashion - an event occurred or it didn't, a defendant is guilty or he isn't, a defending party is deemed civilly liable or they are not.

Another issue with RC, as it relates to our desideratum, is that it seems at odds with our intuitions of when it is reasonable to believe something. Most people recognize that several factors are operative in assessing whether it is reasonable to believe that p. One example is the context - *viz*, what is at stake? Suppose, for example, I am contemplating whether to do in-person banking at my local bank on a Friday at 11 am or the next day at 5 pm. Although I know the bank will be open on Friday morning, I also seem to remember it being open after 4 pm on Saturdays. Despite having never done in-person banking in the evening on a weekend, it might still be reasonable for me to believe that the bank is open after 4 on a Saturday. However, it may also be that being mistaken about the hours of the bank will lead to my personal check bouncing and ultimately being assessed fees. Given what is at stake, it may be unreasonable, then, for me to believe that the bank is open in the late afternoon on a Saturday, even in the absence of countervailing evidence.

The above example lends to the idea that reasonable belief is context sensitive - while, in one context, it may be reasonable to believe p given a body of background evidence; in another context, it may be unreasonable to believe that very proposition,

despite the same body of evidence. For RC to to satisfy the adaptiveness requirement, it must show that it is compatible with the standards of proof in civil proceedings. Furthermore, given that the burden of proof in civil proceedings is preponderance of evidence, this means that RC must also show that it is *always* reasonable for the fact-finder to fully believe the plaintiff's account of events whenever the plaintiff satisfies the burden of proof. However, it is just not true that the fact-finder always thinks it reasonable to fully believe the plaintiff whenever the plaintiff satisfies the burden of proof.

I conclude then that RC may be a sensible theory of punishmnent. However, it does not have the resources to satisfy **DII**.

CHAPTER 4

Epistemology

Let us take stock. So far, I have considered several frameworks as possible candidates to meet the desiderata stated in Chapter Two. As was shown in Chapter Three, however, no framework meets each requirement. Despite their shortcomings, many of the frameworks offer serious attempts at specifying the central difference between individualized evidence and statistical evidence. As I turn to my positive account, I will use some of the strengths of the previous accounts discussed in the previous chapter as a guide for what is needed to satisfy the above desiderata.

My account, in many ways, resembles Moss' account, which is rooted in a relevant alternatives account of knowledge in epistemology. Given the similarities between our respective accounts, it will be significant to highlight places where we differ. Finally, given my aim is to satisfy the desiderata stated in Chapter Two, my focus will be on showing how my proposal does that and less on showing that it handles common cases in naked statistical evidence (e.g., *The Blue Bus* case, *The Prison Riot* case, etc.,). I begin with a brief discussion of epistemic contextualism from which different versions of relevant alternatives accounts of knowledge emerge.

4.1 CONTEXTUALISM IN EPISTEMOLOGY

Epistemic contextualism is fundamentally a view about the way the word "knows" functions. The central idea is that the standards for applying the word "knows" vary

from one context to another.[95] Specifically, "S knows that p (at t)" can be true in one context and false in another - for the same subject S and the same proposition p (and the same time t). For example, in a high-standards context, "S knows that p" may require for its truth that S have a true belief that p and also be in a very strong epistemic position with respect to p. And yet in another context (indeed a low-standards context), the same sentence may require for its truth, only that S satisfy lower epistemic standards, in addition to S's having a true belief that p. What this means is that for the contextualist, one knowledge attributor can truthfully say "S knows that p", while another knowledge attributor, in a different context where higher standards are in place, can truthfully say "S does not know that p", though each attributor is talking about the same S and the same p at the same time.[96]

4.2 RELEVANT ALTERNATIVES

4.2.1 The Machinery of Relevant Alternatives

One family of theories of knowledge that has had a major influence on contextualist approaches is the so-called "Relevant Alternatives Account" (hereafter, RAA)

[95] Put somewhat differently, the truth values of knowledge-ascribing and knowledge-denying sentences (sentences of the form "S knows that p" and "S doesn't know that p" and related variants of such sentences) vary in certain ways in accordance to the context in which they are uttered. What so varies is the epistemic standards that S must meet (or, in the case of a denial of knowledge, come up shy of meeting) in order for that statement to be true.

[96] (See, Greco and E Sosa, 1999)

first proposed by Fred Dretske[97] in the early 1970s and further developed by Gail Stine,[98] and others.[99]

The basic idea behind *RAA* is that knowing some proposition requires that all relevant alternatives to the proposition be ruled out.[100] Given the role that *RAA* serves in my project, it is useful to give a more elaborate sketch of some of the key features of the theory. The next paragraph offers a worked-out example of *RAA*, followed by a discussion of the machinery behind the account.

Suppose S decides to dine at one of her favorite restaurants. Although the food is exquisite, what really brings her back is the service of Pierre, one highly sought after waiter. In S's last visit at the resturant, Pierre provided great recommendations on dishes that she should try for dinner and dessert. Most of all, upon hearing that she was celebrating a birthday, Pierre displayed his mellifluous voice by singing a happy birthday song for her. Although Pierre does not wait on her on the day that she returns to the restaurant, she does believe that she sees a waiter standing near the private room who looks very much like him. The waiter has the same build as Pierre, showed the same attentiveness for his customers as Pierre did for her. He

[97] (See; 1970, 1981)

[98] (See; 1976)

[99] (See Lewis, D. 1996; DeRose, K. 1995)

[100] On this view, it is not necessary, in order to know p, that one be able to exclude *all* the alternatives to p; instead, what is required is merely the ability to eliminate certain relevant alternatives.

even sang happy birthday to his guests just as Pierre did for her. Suppose, then, that, based on S's evidence, she concludes that "Pierre is the gentleman waiting on guests in the private room." Call this, p. Despite S believing p, according to RAA, there are some alternatives to p that she must rule out in order for her belief to count as knowledge. Although the alternatives to p abound, for our purposes, it is only necessary to consider the following minimal set of alternatives to p: S sees another waiter who resembles Pierre but is not related to him (hereafter, r_1); Arthur, Pierre's twin brother, is the waiter whom S sees in the private room (hereafter, r_2); Suppose the restuarant management decides to deceive S by hiring several waiters who dress, speak, and behave just like Pierre. They also take singing lessons to sound just like Pierre. Thus, the waiter whom S sees in the private room is a cleverly disguised Pierre lookalike (hereafter, r_3). I will call the set of alternatives to p as R. Moreover, let us call this example **WAITER**, short for good waiter experience example.

We might represent **WAITER** in terms of a model employing the following variables:

s= the subject to whom knowledge may (or may not) be attributed;
p= the proposition that is believed;
R = the set of alternatives to p; and
e = the subject's ability to eliminate alternatives to p

A pictorial representation of **WAITER** is shown in figure 4.1:

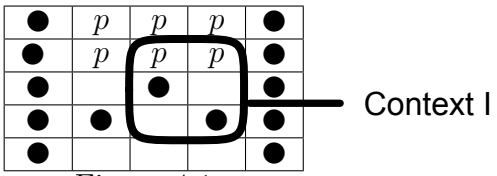

Figure 4.1

To understand the above figure, let each cell on this 5x5 matrix represent a possible world. This means there are twenty five ways in which the world could be. Picture the worlds in which p is true (i.e., p-worlds) as the set of cells containing p, and the alternatives to p as the surrounding cells that do not contain p.[101] The dotted cells represent the set of non-p worlds that s eliminates. Finally, the circle, which is superimposed on the figure, represents the worlds that are in play based on the context.[102] Thus, in Figure 4.1, the worlds being considered are the ones within the circle, which are generated by context I. We'll let the context determine not only the size of the circle, and consequently how many worlds that are being considered, but also the non-p worlds that must be ruled out in order for the agent to possess knowledge. Thus, suppose context I is low-stakes, and selects only the cell in the center of the matrix (i.e., the cell that is in column 3, row 3) as relevant and thus necessary for elimination. Then, s's belief that p counts as 'knowledge' in that context, given that the only relevant cell is eliminated. Of course, other criteria can be used for what constitutes a relevant alternative. If a second context (that is, a higher-stakes context) selects all but one non-p worlds, and context III (an extremely high-stakes context) selects *all* non-p world, as is shown in Figures 4.2 and 4.3 respectively, then s's belief that p does not count as 'knowledge' in either of those contexts.

[101] That is, suppose all the non-$p$worlds are the set of worlds contained in R.

[102] Judgments of the quantity of worlds that fall within the circle and the size of the circle are ultimately determined by the author, and are ultimately exogenous to the set up.

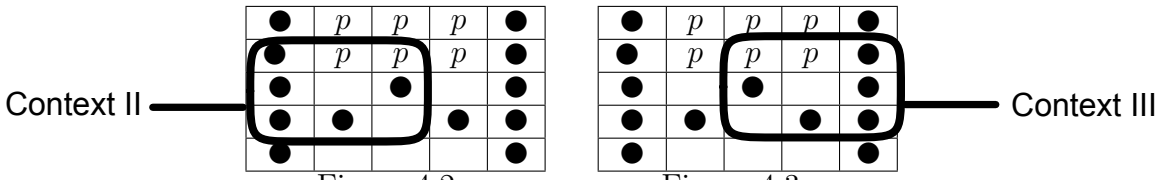

Figure 4.2 Figure 4.3

Schaffer's sketch of *RAA* is useful in illustrating the common features of this theory of knowledge. Despite its virtues, there are still some questions we might have about *RAA*. Perhaps, the most commonly asked question can be put as follows: What makes one alternative, and not another, *relevant*? Understandably, philosophers have provided different answers to this question. Some hold, (a) the relevance depends on facts about what the knower believes or *has evidence* for believing.[103] Others argue that (b) relevance is affected by *facts about the knower's environment*, even if the knower is unaware of those facts. Still others believe that (c) whether some alternative counts as relevant doesn't depend so much on facts about the *knower*; rather, it depends on facts about the *conversational setting in which knowledge is being ascribed.* While I do not wish to settle the issue here, I maintain that something similar to (c) is what is needed for my overall aim. One attempt to lay out a set of rules for (c) is found in Lewis (1996), which I develop in the next Subsection.

4.2.2: Lewis on Relevant Alternatives

David Lewis' own version of *RAA* is set forth as follows:

[103] (See Austin)

S knows that p iff S's evidence eliminates every possibility in which not p - Psst! - except for those possibilities that we [ascribers possibly distinct from S] are properly ignoring.[104]

This means that a subject has propositional knowledge that p only if p is true in every relevant uneliminated possibility. For Lewis, whether or not a possibility is relevant is determined contextually, and he puts forth seven rules of relevance which clarify how context determines relevance. The rules are divided into two sets. The first three principles are commonly called prohibitive rules, which determine the range of possibilities that may not properly be ignored. The remaining four are permissive principles, which determine the range of possibilities that may be properly ignored. Although the rules are fairly straightforward, it is worthwhile to show each rule along with a brief description of it.

RULE OF ACTUALITY: The possibility that actually obtains may not properly be ignored. • "[I]t is the subject's actuality, not the ascriber's, that never can be properly ignored."

• "[The subject's actuality is, for example, the possibility of being the subject on such-and-such a day at such-and-such a possible world.

RULE OF BELIEF:: A possibility may not be properly ignored if the subject gives it, or ought to give it, a degree of belief that is sufficiently high.

• But how high is high *enough?* Suppose that we are in in criminal court, where there is more at stake than there is in civil court. We want to know whether Darren committed the crime. If we give some alternative to that claim - say, that Kofi

[104] (Lewis; 1996, 425)

committed the - even a slight degree of belief, that alternative can become relevant (and we'll need to eliminate it in order to know that Darren committed the crime).

RULE OF RESEMBLANCE: We may not ignore a possibility that saliently resembles a relevant possibility. • This rule is supposed to be the rule that allows Lewis's theory to account for lottery paradoxes[105] and Gettier cases.[106]

RULE OF RELIABILITY: We may ignore a possibility in which a reliable process (such as perception, memory, or testimony) fails.

RULE OF METHOD: We are entitled properly to ignore possible failures in sampling (*induction)* and in inferences from the best explanation (*abduction)*

RULE OF CONSERVATISM: Generally ignored possibilities may properly be ignored.

RULE OF ATTENTION: A possibility not ignored at all is *ipso facto* not properly ignored.

The above rules are regarded as rules of thumb[107] that often determine whether a possibility is relevant in a given case.[108]

[105] For a discussion of the lotter paradox, see Wheeler (2007)

[106] For a discussion of Gettier counterexamples, see Gettier (1963)

[107] (See Schaffer; 2001, 203)

[108] For Lewis, a possibility w is eliminated for S just in case S's perceptual and memory in w would not exactly match his perceptual experience and memory in actuality. For Lewis, perceptual experience and memory are forms of basic evidence. Lewis' general idea is that eliminated possibilities are those in which the subject's basic evidence would differ. So, on Lewis' account, S knows that p iff the rules of

CHAPTER 5

From Epistemology to the Law

Using the combination of Schaffer's diagram of RAA and Lewis' set of rules of relevance (hereafter, the Schaffer-Lewis model), we have a more plausible model for a theory that can satisfy the desiderata stated in Chapter 2. Several features of the joint account, in particular, are instructional for our purpose. To begin, on the Schaffer-Lewis model, knowledge may properly be attributed to a subject just in case she can eliminate every relevant possibility of error.

5.1 Developing *LRA*

Ruling out possible alternatives to a proposition also figures prominently in legal decisions. The legal fact-finder, for example, must verify a set of propositions made by each party by relying on such tools as oral testimony, documents and physicial evidence.[109] This evidence, in conjunction with generalizations derived from personal experience and social convention, help the fact-finder make factual inferences in order to evaluate relevant alternatives to the claimant's narrative of the events that need to be explained.

While the trier of fact is concerned with whether each party speaks truthfully, this

relevance do not select any possibility in which p is false but S's basic evidence stays the same.

[109](See Pattenden; 2009, 3)

is not the former's primary (or, even main) task. Indeed, in order to reach a finding, the fact-finder must also be concerned with the cogency of each party's narrative, which is often evaluated in terms of explaining key events (e.g., whether the defendant did cause the accident). To be sure, the fact-finder makes several interrelated evaluations, including evaluating the possibility of errors that each party's narrative does or does not eliminate; and, evaluating the explananda that each party's narrative has or has not explained.

Obviously, not every error a given party commits (or event they leave unexplained) is weighed equally. Although errors generally should be avoided, some may prove innocuous, and, understandably, are dismissed. Similarly, some unexplained events may seem inconsequential to any of the party's case; while others could prove devastating. How ought the fact-finder to determine how to assess each error or explanandum? The Schaffer-Lewis model also sheds light on how this might be answered within a legal framework. In attributing knowledge to a speaker, Lewis invokes a general set of rules of relevance that may not properly be ignored by the agent to whom knowledge is being attributed. But not each rule of relevance needs to be considered in a given case. Rather, what governs whether the rule is invoked is, *inter alia*, the context in which the speech act is made.

5.1.1: The Machinery of LRA

5.1.1.1: Thesis I

The law seems to function similarly in dealing with the possibility of erroneous claims

or unexplained events. In a given legal proceeding, rules abound for the trier of fact to take into consideration. But, as is true for Lewis' requirement for knowledge attributions, some rules need not be entertained. What determines whether one rule, but not another, is relevant will again be the context, *viz* - the legal proceeding.

The Schaffer-Lewis model has laid the groundwork for, what I shall call, the Legally Relevant Alternatives view (or, *LRA*), which comprises three central claims. The first thesis, hereafter, **I**, states that the varying burderns of proof across legal proceedings can be captured, with a few modifications, by the relevant alterntatives account of what shifts for knowledge attributions.

To illustrate, suppose that Julian is arrested after officers from the Sherrif's Department responded after receiving an emergency call that a man has been injured. Eyewitnesses later tell police that a man had brutally attacked another man sitting in a compact car. Upon arriving at the scene, the officers see Julian driving away and the officers pursue him to an apartment, which the latter uses as residence. Inside the apartment, the officers find a burner phone that contains a caustic text correspondence between Julian and the man in the car. Moreover, during interrogation, Julian admits that the man in the car owes him some money, but denies assaulting him. At Julian's arraignment, he pleads not guilty and thus goes to trial. For convenience, I will call this example **ATTACK**, short for brutal attack example.

We might represent **ATTACK** in terms of a model employing the following variables:

t = trier of fact;[110]

[110] The knowledge analogue to this is the subject (or, s)

f = a finding for the claimaint (in this case, the state);[111]

R_L = the set of legally relevant alternatives to the claimant's account of events[112]; and

$e = t$'s ability to eliminate alternatives to the claimant's account of events[113]

A pictorial representation of **ATTACK** is shown in figure 5.1:

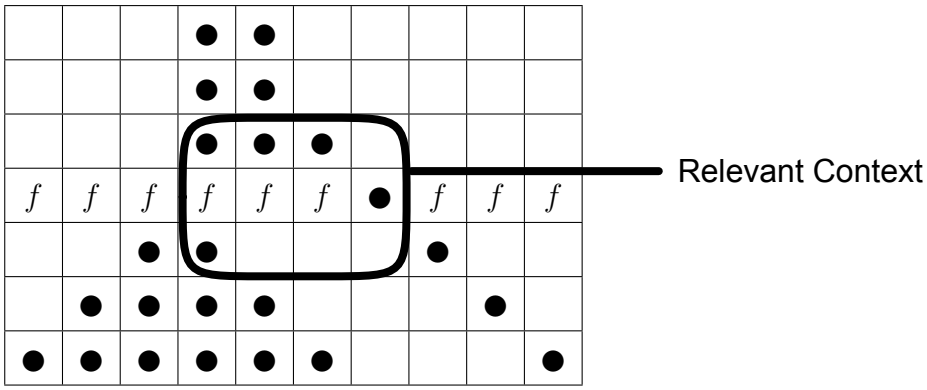

Figure 5.1

As will become apparent, the machinery of RAA easily translates to the machinery I use to undergird LRA. For example, the f worlds are the set of boxes containing f, and the alternatives to f as the surrounding boxes that do not contain f.[114] Also, as was the case for the machinery of RAA, the context - *viz*., the legal proceeding of the case - determines what suces for the relevant alternatives. Finally, the dotted

[111] The knowledge analogue to this p.

[112] The knowledge analogue to this is R

[113] The knowledge analogue to this is e.

[114]That is, suppose all the non-p worlds are the set of worlds contained in R.

boxes represent the set of non-f worlds that t eliminates. For RAA, the elimination of *all* relevant alternatives to p is required for knowing that p. However, for a fact-finder to meet the burden of proof for the relevant proceeding, I do not require the elimination of all relevant alternatives. Instead, it will suffice if the fact-finder eliminates a *sufficiently high* number of relevant alternatives.

5.1.1.2: Thesis II

One lingering question some will have about my account might be put this way: What makes the burdern of proof in one legal proceeding more (or less) stringent than that of another proceeding? LRA has a ready-made answer, which is summed up in thesis **II**: Similar to the role that Lewis' rules serve in considering the range of possibilities to determine whether an agent possesses knowledge, the law also imposes rules that the fact-finder must use to consider the range of possibilities when deliberating whether to find for the plaintiff. Of course, the quantity of rules is case relative: some cases generate more prohibitive rules than others.

I also claim that the varying burdens of proof are, in part, determined by the number of rules - primarily, rules of relevance - that are triggered in the legal proceeding that the case is in. Below is a minimal set of prohibitive rules (hereafter, rules) available to the fact-finder in a given legal proceeding. My goal, here, is not to explain each rule; rather, it is to present a typical list of rules available to the fact-finder and show how some rules are pertinent in some cases but not in others:

(Facts) Facts of the case may not properly be ignored.

(Legal statutes) Laws pertaining to the case may not properly be ignored.

(Evidence) Direct or indirect evidence may not properly be ignored.

(Credibility) Credibility of witnesses may not properly be ignored.

(Inferences) Reasonable conclusions from direct and indirect evidence may not properly be ignored.

(Arguments) Arguments of counsel may not properly be ignored may not properly be ignored

(Special instructions) Instructions of the court may not properly be ignored.

(Impeachment) Evidence of other acts, including prior felonies, by defendant may not properly be ignored

(Silence) Defendant's silence in the face of accusation may not properly be ignored.

Notice that the above rules do not apply in every legal proceeding. For example, while rules 1 - 9, may seem applicable in a homicide case, only 1 - 7 seem to apply in a dispute with your landlord over a security deposit. Similarly, rules 1 - 8 may be applicable in a standard burglary case, but not in a standard tort case. Ultimately, each legal proceeding has its own set of rules that the trier of fact must consider in reaching a decision, and while there is overlap among the proceedings, they usually do not share the same set of rules.

5.1.1.3: Thesis III

Finally, thesis **III** of *LRA* holds that the law prefers types of evidence that tend to yield high degrees of certitude. Indeed, one requirement of admissible evidence is that it tends to prove the matter in dispute. This is why such evidence as video footage, audio recordings, the murder weapon, etc., are admitted in court. When used to support a claim, they tend to raise the probability of that claim being true.

The third thesis is merely an extension of the requirement that admissible evidence be reliable. If the law requires that admitted evidence has the property of raising the probability of a claim under dispute being true, then the more of the property that form of evidence possesses the more the law tends to endorse that form of evidence (or, even the evidence-producing method).

5.2 Understanding the Sufficiency Condition

Above, I state an important difference between *LRA* and the version of *RAA* on which my proposal is based. Specifically, *RAA* requires the elimination of all relevant alternatives in order to satisfy some requirement; my proposal, however, denies that the elimination of all relevant alternatives is needed to satisfy some requirement. It turns out that Moss' probabilstic knowledge account has a similar requirement as that of *RAA* - she requires that the trier of fact eliminates *all* relevant alternatives in order to be in a position of knowing that the standard of proof has been satisfied for a given proceeding. This marks an important difference between *LRA* and Moss' probabilistic account. I think it is too onerous to ask the trier of fact to eliminate *all* relevant alternatives to the plaintiff's narrative in order to find for the plaintiff. Indeed, such a requirement would yield an undesirably *low* number of findings for the plaintiff. This is one of the grounds for my position that a finding for the claimant requires that only a *suciently high* number of relevant alternatives to the claimant's narrative be eliminated.

Despite the advantages to my departure from Moss' view, there are still questions that *LRA* generates. One question concerns the requirement for a suciently high number of relevant alternatives to be eliminated. We can put the question this

way: How does one know when a suciently high number of alternatives have been eliminated? For conveneince, I will refer to this question as the suciency question, or the suciency condition. The rest of the chapter explains how the suciency condition can be met on *LRA*.

5.2.1 Resilience

5.2.1.1 Belief Perseverance

To answer the question of when enough alternatives have been eliminated, I turn to a curious phenomenon people exhibit when confronted with evidence that challenges an occurrent belief: belief perseverance. Belief perseverance (or belief resilience) is the tendency to hold onto a belief even after receiving new information that disconfirms or gainsays the grounds of that belief. There are several explanations for our tendency to hold onto false or discredited beliefs. In a survey article, Lee Ross and Craig Anderson (1982) present an experiment designed to explore "the phenomenon of belief perseverance in the face of evidential discrediting:

[s]ubjects first received continuous false feedback as they performed a novel discrimination task (i.e., distinguishing authentic suicide notes from fictitious ones) ... [Each subject then] received a standard debriefing session in which he learned that his putative outcome had been predetermined and that his feedback had been totally unrelated to actual performance. Before dependent variable measures were introduced, in fact, every subject was led to explicitly acknowledge his understanding of the nature and purpose of the experimental deception.

Following this total discrediting of the original information, the subjects completed a dependent variable questionnaire dealing with [their] performance and abilities. The evidence for postdebriefing impression perseverance was unmistakable ...

On virtually every measure ... the totally discredited initial outcome manipulation produced significant "residual" effects.

A recent series of experiments ... first manipulated and then attempted to undermine subjects' theories about the functional relationship between two measured variables: the adequacy of firefighters' professional performances and their prior scores on a paper and pencil test of risk preference ... [S]uch theories survived the revelations that the cases in question had been totally fictitious and the different subjects had, in fact, received opposite pairings of riskiness scores and job outcomes[115]

On the experiments, Ross and Anderson conclude that from these and other experiments that it is clear that beliefs can survive potent logical or empirical challenges. They can survive and even be bolstered by evidence that most uncommitted observers would agree logically demands some weakening of such beliefs. They can even survive the complete destruction of their original evidential basis.[116]

Commenting on the same experiments, Gilbert Harman (1986) adds:

> In fact, what the debriefing studies show is that people simply do not keep track of the justificatiton relations among their beliefs. They continue to believe things after the evidene for them has been discredited because they do not realize what they are doing. They do not understand that the discredited evidence was the *sole* reason why they believe as they do. They

[115] (See pp. 147 - 9)

[116] *Ibid*

> do not see they would not have been justified in forming tthose beliefs in tthe absence of the new discredited evidence. They do not realize these beliefs have been undermined. it is this, rather than the difficulty of giving up bad habits, that is responsible for belief perseverance.

Ross and Anderson, and (especially) Harman, point out some negative aspects of belief perseverance. It can be irrational to hold onto discredited beliefs, especially if the only ground for accepting the belief is what has been undermined. Moreover, a person who forms a habit of knowingly holding onto beliefs in the midst of conflicting or insucient evidence may gain the reputation of being unreliable.

But the critical comments stated above are not all that could be said about belief perseverance. There are some positive aspects, too. For example, when we at times are willing to keep a discredited belief, our willingness to do so may speak to the quality and quantity of the evidence we have for the belief. Indeed, it could be that the evidential support for believing that p is so strong that evidence that appears to conflict with the support for p does not impact our believing p.

Belief reslience is deeply connected to satisfying the suciency condition. Specifically, meeting the suciency condition depends on the likelihood that the judge or jury would be willing to change their finding if provided additional, particularly, conflicting, evidence. But there are still questions that remain. Arguably, the put as follows: What distinguishes the plausible cases of resiliency with the implausible ones? This question should not be ignored. Indeed, it would be epistemically irresponsible to retain a belief in the midst of contrary information. I maintain that what marks the difference between plausible cases of resiliency and implausible ones

is *weighty evidence*, which is a matter of how substantial the evidence is.[117] As the agent acquires more evidence which bears on the claim, the weight of her evidence increases, which in turn may make a difference to what it is reasonable to believe on its basis. There is an intellectual history of evidentiary weight dating back to Charles Sanders Peirce. To appreciate how resilience can be bolstered by evidentiary weight, it will be worthwhile to consider some proposals of evidentiary weight. I will follow my survey on weight of evidence by considering a couple of proposals on resilience, which incorporate evidentiary weight.

5.2.1.2 A Survey of Evidentiary Weight

5.2.1.2.1 Peirce's view

Although it is debated, several scholars[118] find in Peirce's 1878 essay, "The Probability of Induction" the first clear formulation of the notion of weight of evidence.[119] In the essay, Peirce contrasts a materialistic view of probability, according to which probability statements report facts about the frequency with which one kind even accompanies another kind, with De Morgan's conceptualist view, which holds that

[117](See Kelly; 2008, 2)

[118] (See Carnap, 1962, 554; Popper, 1959, 406; Good, 1985, Ullian, 1995, 96-97; and O'Donnell, 1989, 76)

[119] The informal concept of weight has a long history. I. J. Good notes that, as a metaphor, it can be traced back at least to the Ancient Greek goddess Themis and her 'scales of justice', in which the scales involvee both a balancee and a quantity of evideen

probability statements concern one's subject degree of belief that ought to attach to a proposition. Although Peirce goes with a materialistic conception, he offers arguments on behalf of conceptualism. The heart of Peirce's defense is as follows:

> Belief is certainly something more than a mere feeling; yet there is a feeling of believing, and this feeling does and ought to vary with the chance of the thing believed, as deduced from all the arguments. Any quantity which varies with the chance might, therefore...serve as a thermometer for the proper intensity of belief. Among all such quantities there is one which is peculiarly appropriate....As the chance diminishes the feeling of believing should diminish, until an even chance is reached, where it should completely vanish and not incline either toward or away from the proposition. When the chance becomes less, then a contrary belief should spring up and should increase in intensity as the chance diminishes, and as the chance almost vanishes (which it can never quite do) the contrary belief should tend toward an infinite intensity. Now, there is one quantity, which, more simply than any other, fulfills these conditions; it is the logarithm of the chance.[120]

Here Peirce calls the conceptualist's notion of the degree of belief as a measure of the proper intensity of the feeling of belief.[121] By treating the chance as the

[120] (See W3 293 - 294)

[121] To be sure, Peirce maintains that believing involves more than feelings. He claims that belief has three properties - it is something of which we are aware; it appeases

external fact which generates the belief, Peirce leaves the logarithm of the chance as the appropriate measure of belief intensity. Peirce argues then that "belief ought to be proportional to the *weight of evidence*, in this sense, that two arguments which are entirely independent, in this sense, that two arguments which are entirely independent, neigher weakening nore strengthening each other, ought, when they concur, to produce a belief equal to the sum of the intensities of belief which either would produce separately"[122]

It is in developing this last reason for linking degree of belief and the logarithm of the chance that Peirce introduces the term "weight of evidence." It appears that Peirce treats the metaphors of weighing and of balancing evidence as interchangeable.[123] In fact, given that Peirce confines his discussion to cases in which the arguments in question are independent of one another, his characaterization of weight of evidence draws parallels to what Isaac Levi calls the net weight of evidence: "Take the sum of all feelings of belief which would be produced separately by all the arguments *pro*, subtract from that the similar sum for arguments *con*, and the remainder is the feeling of belief which we ought to have on the whole."[124] The logarithm of the chance is taken to represent the facts which "produce" the belief, and then the feeling of belief ought to remain proportional to this magnitude. Kasser writes: "Just as our the irritation of doubt, and it governs action by establishing habits. (See W3 247 and 263)

[122] (See W3 294)

[123] (See, Kasser; 2015, 638)

[124](See W3, 294)

gustatory system, if functioning properly, ought to issue in results that remain proportional to the log of the stimulus that produced the taste, our cognitive apparatus ought to respond to the log of the evidential stimulus.[125]

Despite offering some novel arguments on behalf of conceptualism, Peirce ultimately rejects the view in favor of another position. Peirce's objections with conceptualism are two-fold: he raises an issue with the intelligibility of an appeal to a principle of indifference. Secondly, he denies the legitimacy of personal probabilities not appropriately grounded in facts. Regarding Peirce's second formulation of weight of evidence, some read it as an anticipation of Keynes' conception of weight of argument, which is what I turn to next.

5.2.1.2.2 Keynes' view

In *A Treatise on Probability*, Keynes provides the following insight:

> As the relevant evidence at our disposal increases, the magnitude of the probability of the argument may either decrease or increase, according as the new knowledge strengthens the unfavourable or the favourable evidence; but something seems to have increased in either case,—we have a more substantial basis upon which to rest our conclusion. I express this by saying that an accession of new evidence increases the weight of an

[125] Incidentally, Turing seems independently to have come up with a measure of evidential weight that deploys odds, logarithms, and the analogy with the decibel system almost exactly as Peirce did.

argument. New evidence will sometimes decrease the probability of an argument, but it will always increase its "weight".

Unlike Peirce, Keynes seems to have gross of evidence, not weight of evidence, in mind. To see this, suppose someone (first inquirer) who has drawn a small number of beans, the vast majority of which are black, has good evidence, good reasons for belief, in one clear sense. We might even suggest that such evidence as she has bears decisively on what she is to expect. The net weight of her evidence might not be large, but it points in one direction; namely, in this case it clearly favors the hypothesis that the bag contains many black beans. Someone (second inquirer) who has drawn a much larger number of beans from the bag (with replacement and stirring) but has drawn an approximately equal number of black and white beans is in a less favorable situation for having an expectation about what color the next bean will be. But her judgment is in some clear sense more settled or more stable than the belief of the other inquirer.[126]

Several points about gross (or Keynesian) weight can be gleaned from this example. First, Keynesian weight lacks directionality or valence; the weight of evidence for H will always be the same as the weight of evidence for ¬H. This also means that an integral feature of Keynesian weight is that it always increases when new relevant evidence is learned. Our second inquirer's evidence is weighty in Keynes'

---

[126] Brian Skyrms holds that increasing the number of trials increases the "resilience" or resistance to change of the probability judgment and Joyce characterizes Keynesian weight in terms of the concentration and stability of credences in the face of changing information.

sense.[127] Keynes is a relationist about probability and so denies that probability judgments based on more evidence are more likely to be true than those grounded on less evidential weight. Finally, probability judgments for Keynes are about the relationship between the evidence and the conclusion, and are, when true, logically true.

5.2.1.2.3 Cohen's view

Accounts of evidential weight have also been developed with the aim of interpreting the legal standards of proof. Laurence Jonathan Cohen, for example, has provided a system of probabilities he calls *Baconian* so as to acknowledge their linkage to Bacon's central work about the elimination of hypotheses on the basis of experimental tests. The title of Cohen's major work on Baconian probability is *The Probable and the Provable* (1977). In this work, Cohen interprets probability to be a means for grading the provability of hypotheses (or generalizations) people assert about matters of interest.[128] Cohen adopts a pluralistic view of probabilistic reasoning - different situations require different intterpretations.[129] In some situations an interpretattion of a Pascalian probability[130] may be acceptable. In other situations,

[127] Such evidence would (on its own) produce no strong feeling of belief according to Peirce's method of balancing reasons.

[128] (See 1977, 13-32)

[129] (See 1977, 5-32; 1989a 81-115)

[130] The probability theory associated with Pascal is better known as enumerative induction.

grading the provability of some hypotheses may require probabilities having very different properties. Specifically, Cohen addresses situations in which the provability of some hypothesis can increase only to the extent that the agent is able to eliminate possible reasons for its invalidity.

To understand Cohen's proposal, consider the following example adapted from (Schum, 1977): Suppose I assert the hypothesis: if a thing is a ϕ, then it is also a φ. A common strategy to support this claim is to show instances of the joint occurrence of ϕ and φ- *viz*., enumerative induction. Following Cohen, eliminative induction requires properly ruling out instances of ϕ's that are not also φ's. Finally, Cohen claims that in eliminative induction the best we can do is to say that the test provides only an *ordinal* measure for grading evidential support. Specifically, we can only say that one hypothesis is better supported than another but not by how much or by what ratio.

5.2.1.3 A Survey of Resilience

5.2.1.3.1 Skyrms on Resilience

Not only is there an intellectual history on evidential weight, but there is also one on resilience. To appreciate the connection between both evidential weight and resilience, it is also worthwhile to consider some views on resilience. Our first discussion is found in Brian Skyrms' book *Causal Necessity*, which attempts to deal with many of the major problems facing philosophers today. The central idea of *Causal Necessity* is that invariance is the key to understanding many of these problems. The

notion of invariance is applied to problems such as randomness, epistemic probabilities, conditionals, confirmation, decision theory, just to name a few. I will focus, however, on Skyrms' discussion on stattistical laws, which will lead to a discussion on resilience.

The first part of Skyrms' book deals with propensities and statistical laws. Skyrms believes that propensities are the probabilities that play a role in statistical laws, and his discussion of statistical laws depends on that claim. Statistical laws tell us that certain systems have a stable probability, which Skyrms calls propensities. Skyrrns also explains what it means for a probability to be stable, and he defines a notion of resiliency which is supposed to capture the idea of invariance and stability. Resiliency is defined as:

> Resiliency of $\Pr(q)$s being α= 1 - $\text{Max}_i|\ \alpha — \Pr_j(q)|$ over $p_1 \ldots\ p_n$ (where the $\Pr_j$s are gotten by conditionalizing on some truth-functional compound of the p_is which is logically consistent with both q and its negation)[131]

In tthis definition the p_is are properties or experimental factors which are considered relevant to the occurrence of q. Resiliency measure the independence of q and these factors. Consequently, resiliency is a measure of stability, independence, and invariance. If the resiliency of a certain proposition is 1, then we know that the proposition is necessary or invariant. Degrees of resiliency less than one correspond to cases of approximate independence or approximate invariance. Thus we can look at the resiliency of a proposition to determine how close we are to the ideal.

[131](See Skyrms, 1982, 11-12)

5.2.1.3.2 Lawlor on Resilience

I now turn to a recent view defended by Krista Lawlor, which incorporates elements of weightiness and resilience. Lawlor begins by distinguishing between two concepts: *opinion* and *conviction*. Opinions (or, credences) are the degree of belief that a person may have in a proposition. Relatedly, convictions are a type of strong opinion with a key property - *viz*, they necessitate the anticipation of constancy of the truth of one's opinion.[132] That is, to have a conviction about p means that one anticipates holding p even under new information. Suppose Suzy, for example, purchases a lottery ticket where her chances of winning are one in a billion. She may have the strong opinion that her ticket is a losing ticket, while still being open to changing her opinion upon new information. If, however, she reads the newspaper and learns that she did not win, she may conclude that the matter is settled, though her opinion only slightly increases. Although Suzy has a strong opinion about her ticket losing before reading that it does, Suzy's opinion is not a conviction, given that she anticipate her belief changing upon new information. It is only when she concludes that the matter is settled that her opinion elevates to a conviction. *Anticipatation of constancy of truth*, then, is what distinguishes opinion from conviction.[133]

Of course, mere convictions don't imply that the agent is always being rational.

[132] (See Lawlor; 2013, 6)

[133] Given that convictions require the anticipation of constancy of the truth of one's opinion, we might think that convictions are second-order attitudes. Lawlor, however, argues that anticipating constancy is a first-order attitude, showing itself in a range of behavioral dispositions. (See; 2013, 7)

One might have an attitude that is resistant to counter-evidence. I may dogmatically believe that my guardians are my biological parents, despite DNA evidence and friends and family members suggesting otherwise. Thus, there is a difference between an attitude that is resistant to counter-evidence and one that thinks it will survive new information. Recognizing this distinction, Lawlor develops a notion of *rational conviction*, which is based on conviction. Rational conviction has two central features. The first is good (or, what Bayesian theorists call *weighty*) evidence. To illustrate, suppose Suzy tosses a coin four times and it comes up heads twice and tails twice. Upon observing the throws, she might be inclined to think that the coin is fair. If the coin, on the other hand, were tossed 10,000 times and comes up heads 5,118 times, that is much more weighty evidence for the proposition that the coin is fair. It is important to note some significant aspects of weighty evidence. It can be psychologically impactful. Lawlor writes: "Having weighty enough evidence puts one in a position to anticipate no further change in the truth of one's opinion."[134] Indeed, having weighty evidence "moves one from opinion to conviction." Seeing the coin flipped a handful of times might incline Suzy to believe that the evidence favors the coin being fair, while still entertain the possibility that more evidence will change my mind. If, however, she sees the coin flipped 10,000 times with 5,118 heads, she would rationally anticipate no rational change in her opinion on further inquiry. Moreover, weighty evidence for conviction has implications on action. Thus, just how weighty the evidence needs to be for conviction depends on the action and the stakes involved. Betting my remaining years of freedom requires more weight of evidence than betting a nickel.

[134](See 2013, 8)

Another feature of rational conviction is the rational anticipation that what one is confident of will survive new information. Lawlor develops this property by adopting Brian Skyrms idea of resiliency that aims to capture the 'sureness' that some of one's first-order probabilities enjoy. To see this, suppose Suzy has not seen a coin flipped yet, but she has no reason to think it biased, and so her degree of credence that the coin will land heads on the next throw is $\frac{1}{2}$. After seeing the coin flipped 10,000 times with 5,118 heads, her degree of credence that the coin will land heads on the next throw is also (roughly) $\frac{1}{2}$. But she is surer about it now. Skyrms suggests that as theorists we do not need second-order probabilities about first-order probabilities to capture this sureness. We can instead think in terms of one's reluctance to change one's credences given more evidence. Put somewhat differently, we can think of how much one's initial credence changes, conditional on additional information. The less change in one's conditional credences, the more resilient one's unconditional credence is. Lawlor adopts Skyrms' notion of resiliency, which is set forth as follows:

> Resiliency of Pr(q)s being $\alpha = 1 - \text{Max}_i|\ \alpha - \text{Pr}_j(\text{q})|$ over $\text{p}_1 \ldots \text{p}_n$ (where the Pr_js are gotten by conditionalizing on some truth functional compound of the p_is which is logically consistent with both q and its negation). [135]

Focusing on the proposition q, assume the p_is as propositions that encode information about how the world is or might be. Put differently, the p_is are epistemically possible propositions or ways the world might be for all one knows. The conditional probability of q given the information that p_i measures one's degree of credence in q

[135]See Skyrms, *Causal Necessity*, 11-12

under the supposition that p_i describes the world. In sum, Lawlor holds that rational belief is driven by evidence and permits action without further inquiry because it is based on sufficiently resilient credence.

Practical stakes help set the threshold

There are still aspects of rational conviction that have not been addressed, including matters like a threshold, indeterminancy, absolute resiliency, etc., Lawlor addresses each of these issues. For instance, Lawlor maintains that conviction is a response to firmness of evidence – how much weight or firmness of evidence is needed for conviction is determined in light of one's practical stakes and interests in the question at hand. She writes:

> It is only if the evidence is firm enough (sufficiently unlikely to shift or be undermined) in light of these interests, that one will be persuaded by it, and form a conviction. Similarly, conviction requires resiliency of opinion, but how much resiliency is required is itself determined in light of one's practical stakes and interests ... Belief involves sufficiently resilient credence, where what is sufficiently resilient may change with features of the context, broadly construed to include what is at stake in being right and wrong about the target proposition.[136]

[136]See (2013, 12)

Indeterminancy

Lawlor's goal is not to offer an algorithm that precisely spits out each case in which a person possess a rational conviction. Indeed, such a goal would require a book unto itself. Thus, some indeterminancy is to be expected, since it may be dicult to say precisely where the resiliency threshold is that turns one's high degree of credence into a belief. She writes:"We do not need to digress far here, but it is worth noting that there may be imprecision in measures of resiliency." This is to be expected, since we may only be able to assign rough or imprecise values to one's credences, and these values in turn determine resiliency.

Absolute resiliency vs. restricted resiliency

Finally, Lawlor addresses the concern that there is always some information that would upset one's conviction. It is always possible, for example, that the prosecutor has suborn perjury and so all the witnesses have given false testimony. Although they seemed very credible, it is still the case that one's credence in losing might go down a bit conditional on this possibility. Should this matter? Should one then not have conviction even now on reading the news? Resiliency is a property of credences, and it is relative to the propositions one conditions on. Absolute resiliency in one sense involves no difference between conditional and unconditional credences. Whatever information one considers, one's credence in q does not shift even a little bit. (For most empirical propositions, this is a tall order.) Absolute resiliency has another sense, involving little change in one's conditional credences in q, over all manner of information. Given resource-bounded agents, the set of propositions to condition

on in order to determine resiliency is a subset of the universe of propositions. This subset includes all and only those propositions that express information one judges it relevant to consider. So there are two aspects of resiliency thresholds: how close the conditional credences are to one's unconditional credence, and how many propositions are conditioned on or considered relevant to determining resiliency.

Summation

Summarizing these points: resiliency for a target proposition q requires that for a set of relevant propositions $\{k_i \ldots k_n\}$ in context C, the conditional credences Pr (q | k_i) are not excessively lower than the unconditional credence Pr (q). We can further imagine that Pr (q | •) is somewhat indeterminate in nature and not numerically precise. These points suggest a revision in our statement of the resiliency threshold for rational belief: Resiliency Threshold (revised): It is epistemically rational for us to believe a proposition only if our degree of credence in it is suciently resilient in the context.

5.2.2 LRA & The Suciency Condition

Using Lawlor's framework, I now have a more plausible model for a theory of resiliency, which will then explain what it means to satisfy the suciency condition. According to Lawlor, having a rational conviction implies that two conditions are met. One condition is that the belief is supported by weighty evidence. This is also one of the requirements to satisfying the suciency condition in each proceedings. Surely, we do not want the evidential support for the plaintiff's narrative to be weak.

Of course, some will ask *How much evidence is needed for the support to be weighty?* I will answer this question in two parts. First, I do not think it is pro-

ductive to look for a quantitative answer. There are several factors[137] involved in a criminal case that are not amenable to saying something like at least 90% of all the total evidence must support the plaintiff's case. Indeed, a numerical answer would engender more questions than it would answer. Second, and more importantly, the answer to the above question is best answered by Lawlor's second condition, according to which, rational conviction requires that the agent believes that the matter is *settled*. That is, usually when we think that a matter has been seettled it is because we think there is enough evidential support to bring the inquiry to a close.

To be clear, stating that the matter is settled does not mean that there is, in principle, *no* additional information that will change the fact-finder's belief that that plaintiff's narrative is correct. That is, stating that the matter has been settled does not imply absolute resiliency. Rather it means that if additional evidence related to the case were added, the fact-finder's belief that the defendant is guilty would still survive. Stating that the matter is settled is another way of stating that the fact-finder's belief is resilient. Lawlor also spells this out as follows:

> An agent is resilient for a target proposition q only if for a set of relevant propositions $\{k_i \ldots k_n\}$ in context C, the conditional credences Pr $(q \mid k_i)$ are not excessively lower than the unconditional credence Pr (q).

The formulation also requires some clarifying, if we want it to make sense of how it

[137] For example, criminal cases differ in terms of categories. Most criminal systems for states across the United States divide their crimes into several different categories depending on the seriousness of the crime. The major categories include infractions, misdemeanors, and felonies.

fits into the overall picture of the sufficiency condition. First, it is important to note that Lawlor mentions a target proposition q. In a criminal proceeding, there can be many target propositions, usually falling into one of two categories - the main and auxillary target propositions. The main target proposition is usually the claim that the defendant is guilty of all the charges that have been levelled against him. It is the major claim that ultimately states the defendant's fate. For example, in a felony case, where the defendant is accused of armed robbery, the main target proposition might be that the defendant is guilty of aggravated robbery.

An auxillary target proposition, on the other hand, can be any target proposition that is used to determine the main target proposition. Thus, to continue with the above example, if the main target proposition is that the defendant is guilty of aggravated robbery, then an auxiliary target proposition might be that the defendant was at the scene of the crime when the event occurred. There is also a dependence relation between the two classes of target propositions. The fact-finder's belief about the main target proposition depends on her belief about the auxiliary target propositions. But the same is not true the other way around.

Lawlor's formulation of resiliency also warrants a discussion of relevant propositions $\{k_i \ldots k_n\}$ in context C. In order to assess whether (and to what degree) an agent is resilient about a target proposition q, the proposition must be evaluated against the background of relevant propositions $\{k_i \ldots k_n\}$ in context C. These relevant propositions should be understood as additional information that could be presented to the agent (or fact-finder) once the agent's (or fact-finder's) mind has been made up. The idea being that the agent is said to be resilient about a target proposition q when her credence about q does not waver much even when additional

information is presented to her. Put somewhat differently, her credence about q survives additional information.

It is important to point out, however, that Lawlor qualifies the type of additional information that may be presented. Resiliency does not require little to no wavering in the face of *any* additional propositions. Rather it requires little to no wavering when being presented with any *relevant* propositions. Thus, there is a natural restriction that is placed on the set of additional propositions that one may add to determine whether her credence changes at all. Surely, there are some scenarios where it would be irrational for the fact-finder's credence not to undergo drastic change. If, for example, the jury were to learn that the prosecution knowingly allowed the use of perjured testimony during the trial and routinely withheld exculpatory evidence during discovery, then such infractions would be grounds for disbelieving the state has met her burden irrespective of what the jury may have heretofore believed. Or, suppose that during the trial, someone other than the defendant credibly and voluntarily confesses to committing all the charges levelled against the defendant. In such a circumstance, it would be appropriate to return a not guilty verdict, even if, prior to that, the fact-finder had a rational conviction that the plaintiff had met her burden. But such scenarios are very rare and are consequently ignored when considering additional information that would change the fact-finder's belief in the target proposition.

The above remarks lead me to making the following proposals:

- **General rule**: The sufficiency condition is met just in case the fact-finder has a rational conviction that the standard of proof has been met for the proceeding the case is in.

- **Preponderance of evidence standard**: Requires the rational conviction that it is more likely than not that the defendant is liable.
- **Clear and convincing evidence standard**: Requires the rational conviction that is very likely that the defendant is liable.[138]
- **Beyond reasonable doubt standard**: Requires the rational conviction that it is nearly certain that the defendant is guilty. [139]

5.3 Is conviction just as onerous as knowledge?

Now that *LRA* has been presented, it should also be clear that it differs from other proposals stated in Chapter Three. One proposal, in particular, that I have attempted to distinguish my view from is Moss' probabilistic knowledge account of legal proof. Although my and her accounts rely on a relevant alternatives account of knowledge from the contextualism literature, there are important aspects in which my account differs from hers. Perhaps the biggest difference centers on how we show what is required to satisfy the burden of proof. On her view, the fact-finder must know that the burden of proof has been satisfied in criminal proceedings, and she must know that the burden of proof has likely been satisfied in civil proceedings. As

---

[138] Put differently, the fact-finder must have the rational conviction that the defendant's liability is greater than .7.

[139] Put differently, the fact-finder must have the rational conviction that the defendant's guilt is at least .95.

stated previously, my view denies that knowledge is required to satisfy the burden of proof in any legal proceeding. Although I provide several arguments for rejecting knowledge-based accounts of legal proof, in this section I focus on one particular reason and how it relates to my requirement of rational conviction.

One reason for denying that knowledge is required for satisfying the burden of proof is that it would be overly demanding to require that the fact-finder knows that the burden of proof has been met. To my demandingness objection, Moss, and other proponents of knowledge-based accounts, may reply that knowledege accounts of legal proof are not the only ones that are demanding. Given that I require, in at least two proceedings, that the fact-finder has a rational conviction that the plaintiff has satisfied the burden of proof, this in some respects may be seen as demanding, indeed overly demanding.

My response to this objection is two-fold. First, I believe that I have offered a case for why it is reasonable for the fact-finder to have a rational conviction that the plaintiff has satisfied the burden of proof. I leave it now up to the reader to decide whether the case is a good one. In what I follows, I wish to point out why the demandingness objection that I raise against knowledge-based accounts are starkly different from the demandingness objection that may be levelled against my proposal. The chief ground for why I maintain that it is overly onerous to demand that the fact-finder knows that the plaintiff has met her burden is that I believe that it is quite possible for the burden of proof to be met in cases of wrongful conviction. It is quite possible for the plaintiff to demonstratively show that the defendant is at fault even if turns out that they are not at fault. My account, on the other hand, does not have this problem. By requiring that the fact-finder has a rational conviction that

the burden of proof has been satisfied, that does not imply that the defendant is in fact at fault. Thus, whatever one might mean by stating that *LRA* is demanding, it is not akin to the way in which knowledge-based accounts are overly demanding.

CHAPTER 6

LRA and the Desiderata

6.1 *LRA* and DNA evidence

In this Chapter, I return to the desiderata outlined in Chapter 2. I begin by showing that *LRA* respects **DI**, which requires an explanation for why DNA evidence is treated differently in contrast to other forms of statistical evidence. By the end of the discussion, it will become clear that the explanation that *LRA* provides for the exceptional treatment of DNA evidence matches intuitions on how we view DNA evidence.

6.1.1 DNA worked out

Recall in ***DNA*** that there is very strong DNA evidence supporting the claim that *Defendant* committed the crime of which he is being accused. ***LRA*** can capture why we think that *Defendant* is guilty of the charges. *LRA* holds that a finding for the defendant implies that a sufficiently high number of relevant alternatives have been eliminated. Furthermore, the sufficiency condition is satisfied only if the fact-finder has a rational conviction where having a rational conviction means (i) the fact-finder's belief is supported by weighty evidence; and (ii) the fact-finder's belief that the plaintiff met her burden is resilient.

I believe that the facts associated with DNA evidence typically leads to our having a rational conviction in most claims established by DNA evidence. To see this, note that when a person's DNA profile is said to match the DNA sample collected at

a crime scene at every locus tested, we might still want to evaluate the probability of finding the DNA profile in someone else. That person is assumed to be a random member of the population of possible suspects. In order to rule out a random person, forensic scientists calculate the frequency of the profile in the most relevant population. The frequency is bettter known as the random-match probability, which is an estimate of the answer to the question: What is the probability that a person other than the suspect, randomly selected from the population, will have this profile? Forensic scientists believe that the smaller that probability, the greater the likelihood that the two DNA samples came from the same person.[140] Indeed, the calculated probability of a match between suspect and DNA evidence is usually one in many millions or billions. [141]

By peering further into the random match probability, several connections between it and the requirements for rational conviction can be made. Recall that the smaller the probability the greater the likelihood that the two DNA samples came from the same person. But another point might be made here: the smaller the probability the more possible worlds are eliminated. This should be easy to see. As I decrease the

[140] Put somewhat differently, if the probability is very small, then either the two samples came from the same person or a very unlikely coincidence has occurred.

[141] For example in *Knight v. State* (1993), the expert stated that the probability of selecting an unrelated individual of the population from the same race as Knight who had a genetic profile matching the semen taken from the victim was one in ten billion. Also, in *Dubose v. State*, the jury was told that there is a 1 in 500 million probability that the appellant's particular DNA pattern would appear in the North American black population.

probability that the sample came from the someone other than suspect, I am also eliminating many, indeed, very many relevant possible worlds.

Moreover, the elimination of possible worlds is tied to weighty evidence. As one's overall evidence grows in favoring one side over the other, this also means the elimination of possible worlds that are incompatible with the targeted claim. Given the number of individuals that are generally ruled out in DNA cold-hit cases, this presumably explains why courts have looked upond such cases favorably. The more possible worlds that are ruled out the higher degree of certitude is achieved in concluding the person whom the two DNA samples match is the perpetrator. Indeed, courts presumably regard DNA cold-hit cases with astronomical match probabilities as having an *in-built* weightiness attached to them.

A third connection can be stated. As stated in chapter 5, there is a strong connection between weighty evidence and resilience. Specifically, the weightier the evidence the more resilient one becomes, even if given additional information. Thus, the more astronomically low the match probability is, the more possible worlds are eliminated. Furthermore, the more resilient I am likely to be in concluding that the suspect is the perpetrator. Finally, weightiness and resilience are the two essential ingredients for rational conviction. Thus, an astronomically low RMP can in turn generate rational conviction in the fact-finder, provided the probability is supported by weighy evidence that produces resilience.

With this in mind, we can easily predict how *LRA* would handle ***DNA***, where the match made between the DNA profile recovered from the victim's neck and a DNA databank of offenders supports a 1 in 10^{10} chance (or, an extremely low probability) that someone other than *Defendant*'s DNA was found on the victim's neck. Most

of us would believe that there is ample evidence to conclude that the DNA evidence collected from the victim's neck and the DNA profile match on CODIS belong to one and the same individual. Furthermore, our belief is resilient. That is, we think that questions regarding whether the DNA evidence found on the victim's neck actually belonging to *Defendant* have been settled. Were additional information added, it is very unlikely that our belief would change.[142]

6.1.2 Brawl worked out

While the fact-finder's belief that *Defendant* committed the crime in ***DNA*** satisfies the conditions for being called a conviction, the same cannot be said about the fact-finder's belief in the guilt of the defendant in *Brawl.* At first glance, in ***Brawl***, we might be tempted to think that a greater than .95 chance that *Defendant* is guilty is fairly high. However, if we look at the numbers used to generate that probability, what we see is that by randomly selecting one of the protesters there is still a 1 in 21 chance that an innocent person is wrongly punished. This fraction is not weighty. Under no inquiry would we conclude that the numbers will generate the

[142] Historically, DNA evidence carries a really strong and persuasive power in Court because the random match probability (RMP), defined to be the probability that a person picked at random has the same DNA profile as the evidentiary sample, is very low if several unlinked loci are typed. For a a 13-locus CODIS profile, typical RMPs are on the order of 10^{-14} to 10^{-15}. Such a low RMP implies that a particular DNA profile has a high probability of being unique. Of course, the method of using (RMP) to secure a conviction is just a roundabout way of displaying relevant alternatives being eliminated. By eliminating enough alternatives to a hypothesis, the uneliminated hypothesis left standing is often seen as unique - that is, the correct one.

kind of belief resiliency that is needed to hold onto a belief in the midst of additional (and, possibly, conflicting) information. We may grant that the fact-finder partially believes that *Defendant* took part in the brawl. Indeed, we may even grant that the fact-finder's belief is (partially) reslient. But it is not nearly as resilient as it would need to be to satisfy the suciency condition.

Before moving on to how *LRA* handles **DII**, it is worth emphasizing a point that is easily overlooked. *LRA* not only gives the right rulings in ***DNA*** and ***Brawl***, but it also captures our reasoning in both cases. Consider, for example, ***Brawl***. Many of us think that the court should not return a guilty verdict against the randomly selected protester. Isn't the reason why we think this, because we think, *inter alia*, that the evidence supporting the belief that *Defendant* is guilty is not very weighty? If so, then this seems to be what *LRA* means by concluding that a suciently high number of legally relevant alternatives to the protester actually being guilty would not have been eliminated were punishment to be meted out on him.

Similar points can be made about how *LRA* handles ***DNA.*** We seem to think that the court should return a guilty verdict against *Defendant.* Perhaps our reasoning is that the chances that someone other than *Defendant* committing the crime is suciently low. Of course, if this represents our reasoning, then that is just another way of saying that a suciently high number of relevant alternatives to *Defendant* being guilty have been eliminated.

6.1.3 Possible Objection

No theory is without its critics, and *LRA* is no exception. In this Subsection, I will address one possible objection to my proposal. In Section 6.1, I argue that the

chances of a pro-claimant finding in ***DNA*** will come down to whether a sufficiently high number of relevant alternatives to the claimaint's story have been eliminated. If the elimination of enough relevant alternatives suffices for a finding for the claimant, we might wonder, then, whether something similar can be said of *Brawl*-like cases (i.e., non-DNA cases that rely only on probabilistic evidence). In the original ***Brawl***, it is unlikely that the fact-finder will return a guilty verdict, given that not a sufficiently high amount of legally relevant alternatives to the accused being guilty are eliminated. But, suppose the set up had been different. Suppose instead of twenty one protesters (out of which one is randomly selected), there are five hundred thousand (or, even five million) all but one of whom is responsible for a disturbance that culminates in a member of the legal force suffereing grave injuries. If one person is randomly selected, then there would be a 1 in 5 x 10^5 chance that an innocent person would be menaced were we to mete out punishment. Given the extremely low chance of an innocent man being punished, wouldn't *LRA* be committed, as the objection goes, to conclude that the system of randomly selecting one member from that population can generate a result that satisfies the reasonable doubt burden? Moreover, wouldn't we be obliged to punish the randomly selected protester?

Several thoughts are in order. To begin, the above objection provides an example that is a natural consequence of my proposal. I accept the consequence for a couple of reasons. First, if a non-prejudicial mechanism routinely yields a very high degree of certitude for the matter in dispute, then it should not be immediately discarded. Indeed, it would prove as reliable as DNA evidence. The issue, of course, with non-DNA forms of naked statistical evidence is that they do not routinely yield the same degree of certitude. If they did, we might expect to see attitudinal changes with

respect to cases relying only on probabilistic evidence.

My second reason for standing firm on the afore-mentioned consequence comes from the growing number of court cases suggesting that courts are starting to seem sympathetic to my position. First, some courts have expressed a willingness to allow probabilistic evidence when it assists with the establishment of a useful fact. In *State v Pankow*, a Wisconsin appeals court allowed a testimony by a statistician who stated that the probability that 3 infants out of 20 in the same household would die of sudden infant death syndrome (SIDS) over a 5-year period was "one thousand times smaller than 9.1 in one trillion" (p. 918). The court reasoned that such testimony was permissible because its "sole purpose was to meet Pankow's defense theory that the deaths were attributable to SIDS".[143]

Courts have also shown a willingness to admit statistical evidence when it rebuts a claim. In *State v. Briggs* (1989), a Washington appeals court ruled that it was "obviously useful" for the jury to hear a physician's testimony as to the probability that a stutter in certain situations to rebut the defendant's argument that the absence of testimony than an assailant stuttered exonerated the stutterer-defendant.[144] Finally, some courts are now open to considering the use of statistical base rate evidence when such evidence helps resolve the ultimate issue.[145] In *U.S. v. Rogers* (1985), a federal appellate court ruled that it was not wrong to allow testimony that only 2

[143] (See Kagehiro, Laufer; 1992, 181)

[144] *Ibid*

[145] *Ibid*

out of 1,800 bank robberies in the Los Angeles area involved a person wearing a bandanna: "The fact that very few robberies involve this garb makes it more likely that the same person committed both robberies."[146] Rulings such as the above, coupled with a reduction of restrictions on the use of expert witnesses and scientific testimony (Feinberg, 1989), suggest that reluctance to naked statistical evidence may be on the wane.

6.2 *LRA* and the Adaptiveness Requirement

Another advantage to *LRA* is that it is able to meet the adaptive requirement in a straightforward way. Given that a finding for the claimant requires that a *sufficiently high* number of relevant alternatives to the claimant's narrative be eliminated, this requirement is amenable to the various standards of proof used in different proceedings. In cases where the standard of proof requires preponderance of evidence, the number of eliminated legally relevant alternatives needed to find for the plaintiff will be less than the number of eliminated legally relevant alternatives needed to find for the plaintiff where the standard of proof is clear and convincing evidence. Finally, neither of those standards will require the elimination of as many legally relevant alternatives as that of the standard of proof beyond reasonable doubt.

Recall, in Section 3, Moss' probabilistic account of knowledge was the only framework that meets the adaptive requirement. Similar to my approach, Moss also employs a relevant alternatives account of knowledge as the basis of her view. Yet, on her approach, a standard of proof is met when the fact-finder knows that the set of

[146] *Ibid*

probability spaces according to which the probability of guilt or liability meets or exceeds the relevant threshold for the standard. Simply put, in criminal and civil proceedings, a finding for the plaintiff implies that the fact-finder has knowledge that the legal burden is met. I, on the other hand, do not require that the fact-finder have knowledge that the burden of proof. Having a strong belief, without knowing, that the standard of proof via eliminating enough relevant alternatives can be enough.

6.3 *LRA* and Wrongful Convictions

6.3.1 A Worked-out Example

Finally, *LRA* respects **DIII**. Indeed, this is one consequence of not requiring that the fact-finder know that the standard of proof is met. To see this, consider the following example. Suppose that Billy is arrested and tried for shoplifitng. The prosecution produces video footage that shows him stealing over $4,000.00 worth of merchandise. Additionally, the police find all the stolen merchandise at his home. Billy claims that he is innocent but his alibi does not check out. More importantly, he is not able to explain how the merchandise ended up in his home. At trial, it only takes the jury an hour of deliberation to return a guilty verdict. A year after the trial, the prosecution learns that Billy has an identical twin, Bishop, who was separated from Billy at birth. Although Billy does not know about Bishop, the latter has been studying and following the former. Suppose Bishop has become a career criminal. In his latest scheme, he contrives a plan to make his brother suffer, since their biological mother chose Billy over him. The plan involves Bishop assuming the identity of his twin brother, stealing the merchandise, and planting the items at Billy's home in

order to make it appear that Billy shoplifted. Suppose that law enforcement was not negligent in failing to know about Bishop. They checked out every lead and did their due diligence. In fact, the only reason they learned about Bishop and his devious plan is that the latter felt guilt ridden and confessed to law enforcement what he did. We now have a case of wrongful conviction, where the defendant is found guilty on the basis of very strong but ultimately misleading evidence.

A pictorial representation of this case on LRA is provided in Figure 6.3. Assume we are in a scenario in which satisfying the burden of proof requires that only 50 worlds are eliminated. According to the number of dotted boxes, the standard is met. The jury finds for the plaintiff and is therefore in one of the f worlds. Nevertheless, given what we know, actuality is one of the non-f worlds that has not been eliminated. Thus, we have a case in which the burden of proof is met and yet the accused person is innocent.

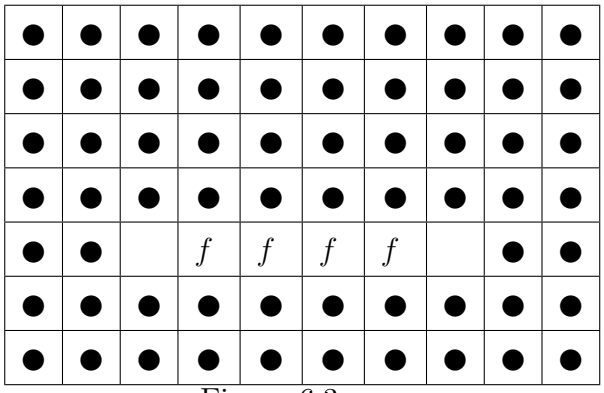

Figure 6.3

6.3.2 The Elusiveness of Reasonable Doubt

Above I show that, on *LRA*, there can be at times wrongful convictions. The feature that allows my account to countenance wrongful convictions is the denial that meeting the burden of proof implies that the fact-finder *knows* that the burden of proof has been met. But, does my rejection of a knowledge-based account come at costs? If so, what are they? In this Subsection, I examine a putative advantage Moss believes her account possesses as a result of requiring that that the fact-finder knows the burden of proof is met in order to find for the plaintiff.

Moss argues that one attractive feature of her probabilistic-knowledge account of legal proof is that it can capture a mystery in the law. The mystery has to do with defining the reasonable doubt standard. Some judges have unsuccessfully tried to develop a definition of reasonable doubt.[147] While others have been strongly opposed to the project.[148] The upshot is that we are left with a standard of proof that no one seems to have a very good handle on.

Despite the difficulty with defining reasonable doubt, Moss believes there is a deeper reason for it being difficult to explain. She then introduces a concept that is also elusive: knowledge. As stated above, Moss adopts a relevant alternatives

[147] See (Whitman; 2008, 2)

[148] To quote a representative opinion from the Oklahoma Court of Criminal Appeals, "We are at a loss to understand why trial courts in this jurisdiction continue to give such an instruction when we have condemned them from territorial days to the present ... it is error for the trial judge to try to define reasonable doubt." (See, *Jones v. Oklahoma,* 554 P.2d 830, 835 (1976)

account of knowledge. On this view, an agent knows p only if she eliminates all relevant possibilities to p. One attractive feature of this theory of knowledge is that it provides a manageable way of making knowledge ascriptions. The agent does not need to eliminate all alternatives to p, only the relevant ones. However, as Moss points out, a relevant alternatives account of knowledge can also be elusive. The reason for this is that whenever one considers an alternative to p, no matter how recherche, it thereby becomes relevant. Thus, even the exotic alternatives that do not need to be ruled out get consideration once we give attention to them.

Moss then builds upon the connection between the elusiveness of knowledge and the elusiveness of the reasonable doubt burden of proof by suggesting that the former best explains the latter. She writes: "The more we say in an effort to spell out the difference between reasonable and unreasonable doubts, the more we call attention to the possibilities that jurors *shouldn't* be considering".[149] Several points seem to be in order about the reasonable doubt standard being explained by a relevant alternatives account of knowledge. First, Moss' point has some initial appeal. *Prima facie*, it seems plausible to hold that if two things share a common property, then perhaps the connection is best understood in terms of one of the items being part of the explanation of the other. Second, it should also be noted that Moss' ground for holding that the elusiveness of knowledge explains why it is so difficult to explain reasonable doubt is not based on mere coincidence. She supports her claim with arguments. Third, given the connection knowledge has with the reasonable doubt burden of proof, Moss then claims that to satisfy the latter means that the fact-finder knows that the burden of proof has been met.

[149]See (Moss; 2019, 7)

In denying that knowledge is required for an account of legal proof, we might, then, think that my account is unable to explain the elusiveness of defining the reasonable doubt standard of proof. However, I do not believe that this follows. In the rest of this Subsection, I aim to show how *LRA* accounts for the elusiveness of defining reasonable doubt.

Recall that on *LRA* a finding for the claimant requires that a *sufficiently high* number of legally relevant alternatives to the claimant's narrative be eliminated. I further claimed that one way to measure whether enough alternatives have been ruled out is by determining the chances that an innocent person would be wrongly punished by relying only on statistical evidence. But I refrained from assigning an ideal threshold. One reason for my reluctance is that sometimes there are other factors unrelated to the elimination of legally relevant alternatives that get factored into the fact-finder's deliberation. These factors include, but are not limited to, matters that are either extraneous to the case or explicitly prohibited by the judge. Consider the following case, which I will call **PRIVILEGE**, which is short for broken doctor-patient privilege:

PRIVILEGE: Suppose Stefano is facing serious charges of crimes he actually did not commit. During trial, Antonio, another person of interest, divulges to Sally, his psychiatrist, in her office and during a counseling session, that he actually committed the crime that Stefano is being charged with. Despite the confession to the shrink being protected under doctor-patient privilege, Sally discloses the confession to a local newspaper, which then disseminates the confession to the public. The jury also learns of the confession. When the judge learns of the psychiatrist's misconduct, he forbids the jury for using Antonio's confession as a basis for acquitting Stefano, given

that the psychiatrist broke the terms of confidentiality.

Although this is a somewhat unusual case, hopefully you begin to see my point. Despite the judge's direct instruction to the jury to disregard the confession, it is very unlikely each individual juror will be able to. Indeed, it seems more likely that most jurors will require more evidence than what is usually reasonable to convict Stefano. This is because something that is prohibited has become a factor to consider, even if it is considered implicitly.

My point in this example could be generalized: It seems that for any case there are factors, including irrelevant ones, that may impede each juror from properly deliberating only the merits of the case. These are among the reasons why I hold that, in providing a descriptive account of how the fact-finder determines whether or not to find for the plaintiff, it is nearly impossible to assign a threshold, which, when met, suffices for a finding for the plaintiff. Thus, the way *LRA* handles the mystery of defining the reasonable doubt burden of proof is by calling attention to what likely happens in many cases; namely, other factors impertinent to eliminating relevant alternatives tend to crowd the juror's deliberation, which then makes it nearly impossible to define what is and isn't reasonable doubt.

6.4 *LRA* and Codicils

6.4.1 The Plan

The preceding sections of Chapter 6 show that *LRA* properly respects **DI** - **DIII**. But perhaps the most challenging desideratum is **DIV**, which seeks to determine whether there is independent justification for the satisfaction of each desideratum on

LRA. Put somewhat differently, is *LRA*'s satisfaction of each desideratum only by specially made features - that is, epicycles - designed only to satisfy the requirement? To answer this question, I will proceed as follows: First, I briefly review how *LRA* handles each desideratum. Then, in the process of each discussion, I consider whether there are independent reasons for that feature existing on *LRA*. If no independent reason can be provided, then obviously *LRA*'s satisfaction of that desideratum is *ad hoc*, and consequently, *LRA* would not satisfy **DIV**. If, however, there is an independent reason for each feature on *LRA* that satisfies each desideratum, then *LRA* respects the last desideratum.

6.4.2 Codicils and DNA evidence

In Subsection 6.1, I show that *LRA* can account for the court's special treatment of DNA evidence. The explanation is that DNA profiling systematically eliminates a sufficiently high amount of relevant alternatives. When there is a match between a DNA sample collected at a crime scene and a DNA profile in CODIS, we have no trouble concluding that the sample and the profile belong to one and the same individual on the grounds that the profiling results in the elimination of *enough* relevant alternatives. It turns, out however, that if the need to explain the special treatment of DNA in the law were not a desideratum, *LRA* would still have the feature that the elimination of a sufficiently high number of relevant alternatives is central to satisfying a burden of proof. One reason that the feature would remain is that it seems to capture *all* cases where naked statistics is sufficient for liability, not just those involving DNA profiling.

Consider, for example, *Kaminsky v. Hertz Corp.* (1979), where a Michigan Appel-

late court ruled that stipulated base rate evidence, which indicates that Hertz owns 90% of all yellow vehicles bearing the Hertz logo "establishes a prima facie showing of ownership or control sucient to prevent a summary judgment [for the defendant]." Similarly, in *Sindell v. Abbott Laboratories* (1980), a plaintiff argued that multiple manufacturers were selling defective products and that one of those products harmed her. Plaintiff, however, could not prove which manufacturer in particular harmed her. The California Supreme Court ruled that in such cases, liability may be imposed to manufacturers in proportion with their share of the market for the defective product.

The foregoing cases show the rare instances in which courts find for the plaintiff on the basis of naked statistics. Notice, however, that *LRA* can easily capture why, in such cases, naked statistics suffice for a pro-plaintiff liability finding: the fact-finder determined that a suciently high number of relevant alternatives to the plaintiff's narrative were ruled out. Thus, *LRA*'s satisfaction of **DI** is not *ad hoc*.

6.4.3 Codicils and the adaptive requirement

It should be apparent that *LRA*'s satisfaction of **DII** is not *ad hoc* either. Recall, that what the plaintiff must prove in criminal proceedings is much stronger than what she must prove in civil proceedings. *LRA* can capture this too. Although the requirement in each proceeding is to eliminate a sucently high number of relevant alternatives, what counts as sucient is relative to the proceeding that the case is in. Thus, in a balance of probability proceeding, the number of eliminated relevant alternatives sufficient for a pro-plaintiff finding are not nearly as high as the number of eliminated relevant alternatives needed for a guilty verdict in criminal court. Thus,

it is the flexibiity of the *LRA* machinery that enables it to satisfy **DII.**

But the machinery of *LRA* is not limited to just capturing the various burdens of proof in the law. In this chapter alone, consider the range of cases that are captured by the flexibility of *LRA*. For example, in Subsection 6.4.2, I show how the *LRA* machinery captures two cases in which naked statistics suce for liability. Moreover, in Subsection 6.1, I show how the machinery captures our intuitive judgments in *Brawl* and *DNA* - two naked statistics cases that most people think should result in different findings. Thus, the flexibility of *LRA* serves many purposes than just satisfying DII**.**

6.4.4 Codicils and wrongful conviction

Finally, *LRA*'s satisfaction of **DIII,** as will become clear**,** does not resort to employing codicils. Wrongful convictions are possible on *LRA*, given that the fact-finder need not eliminate all relevant alternatives to meet the burden of proof in a given proceeding. Thus, so long as a suciently high number of alternatives have been eliminated, the fact-finder may return a guilty verdict. But such a requirement can occasionally lead to the conviction of an innocent defendant.

It should be noted, however, that the denial that all relevant alternatives be eliminated for legal factfinding is just part of the story. Recall that *LRA* is based on the relevant alternatives account of knowledge, according to which knowing some proposition p implies the elimination of all relevant alternatives to p. Thus, by denying that the fact-finder needs to eliminate all relevant alternatives to satisfy a standard of proof, *LRA* is ultimately denying that the fact-finder needs to know that all relevant alternatives are eliminated. This further means that it is the denial

that knowledge is necessary for legal fact finding that allows there to be wrongful convictions on *LRA*.

But are there other reasons for *LRA* to deny that knowledge is necessary for legal fact finding? I believe the answer is *yes*. As stated in Chapter 5, the central reason for *LRA*'s rejection of the fact-finder knowing that the burden of proof is met is not to allow cases of wrongful convictions. Rather, it is to accurately reflect that the law does not require the fact-finder to have knowledge either. Indeed, were the law to require knowledge, then this would have detrimental effects in the law. Many cases in which the trier of fact would have found for the plaintiff would likely have resulted in fewer findings for the plaintiff. Thus, *LRA*'s satisfaction of **DIII** is not *ad hoc*.

CHAPTER 7

The Gatecrasher Paradox

The foregoing chapters show some strengths of *LRA*, particularly in how it respects several desiderata. Of course, there are many other desirable properties we might want a theory of individualized evidence to possess besides those I state in Chapter Two. In the philosophy of law literature, for example, discussions about individualized evidence often focus on solving legal puzzles. Given the interest in such puzzles then, we might wonder how my proposal fares with some legal puzzles that are generated by, among other things, probabilities. In this chapter, I discuss how *LRA* handles one such legal puzzle: *The Gatecrasher Paradox* (also known as *The Gatecrasher case*).

7.1 Understanding the Gatecrasher

The *Gatecrasher case* was introduced by Laurence Jonathan Cohen (1977) and has been debated in the literature with some minor variations. For convenience, I will use the following version:

Consider, for example, a case in which it is common ground that 499 people paid for admission to a rodeo, and that 1,000 are counted on the seats, of whom A is one. Suppose no tickets were issued and there can be no testimony as to whether A paid for admission or climbed over the fence. So by any plausible criterion of mathematical probability there is a .501 probability, on the admitted facts, that he did not pay. The mathematical theory would apparently imply that in such circumstances

the rodeo organizers are entitled to judgement against A for the admission-money, since the balance of probability (and also the difference between prior and posterior probabilities) would lie in their favour. But it seems manifestly unjust that A should lose his case when there is an agreed mathematical probability of as high as .499 that he in fact paid for admission.

Indeed, if the organizers were really entitled to judgement against A, they would presumably be equally entitled to judgement against each person in the same situation as A. So they might conceivably be entitled to recover 1,000 admission-moneys, when it was admitted that 499 had actually been paid. The absurd injustice of this suces to show that there is something wrong somewhere. But where? The paradox poses a conundrum. On the one hand, from a balance of probability standpoint, there seems to be grounds for each rodeo attendee being held liable. On the other hand, most people, Cohen included, believe that it would be wrong to impose liability on a randomly selected attendee based on naked statistical evidence. If the plaintiff would be entitled to receive damages from attendee A, then the plaintiff would be entitled to seek damages from every other attendee, based on the same naked statistics. By parity of reasoning, plaintiff would be legally entitled to force all attendees who did not crash the gate to pay twice.[150]

Since Cohen penned his famous example, scores of papers and commentaries have been written in attempts to resolve the *Gatecrasher's Paradox*. Attempts range from discounting the statistical evidence at play to proposing an additional legal

[150] As Nunn writes, "Once the 502nd defendant has been found liable for trespass, there is a certainty that at least one innocent rodeo attendee has been found liable."

rule which would require that evidence be tailrord to a single defendant.[151] Before discussing how *LRA* addresses the problem that the *Gatecrasher Paradox* presents, it is worthwhile to discuss historical responses detailing the various approaches and legal doctrines that commentators have grappled with.

7.2 Historical Solutions

7.2.1 Keynesian evidential weight

In Chapter Five, I discuss the concept of Keynesian weight in order to develop the notion of resilience that makes up the suciency condition. It turns out that Keynesian weight (with some variations in terminology) is also tied to the reason some believe that liability should not be imposed on a random spectator in the *Gatecrasher*. The idea can be found in Cohen (1977), David Kaye (1979), Richard Posner (1999) and Alex Stein (2005). The basic idea, according to this view, is that there is something dubious in the *Gatecrasher Case*: Why is there no other evidence than naked statistical evidence?[152] Why has the organizer of the rodeo not produced

[151] Here is a more comprehensive account. Some evaluate the paradox under, what I will call, a reconceptualized burden of proof: (See Chen; 2013, 1269 -71) Others defend the plaintiff (See Eggleston; 1980, 678). Still others discuss the paradox in conjunction with the Blue Bus problem (See Fienberg; 1986, 693 - 98). Others discount the statistical evidence at play in the Gatecrasher's Paradox (See Kaye; 1979, 36-39) Some attempt to resolve the paradox by proposing an additional legal rule which would require that evidence be tailored to a single defendant, rather than allowing evidence to remain applicable to an entire population of possible defendants. (See Williams; 1979, 297 - 301)

[152] This type of inquiry can also be applied to the Prison Riot Case and the Blue

any evidence that the defendant did not pay for admission, other than the brute fact that the defendant was in attendance at the rodeo? Have serious attempts to secure more evidence been made? The absence of evidence seems to suggest that the investigations are sloppy. According to Posner, it is this suspicion that gives us qualms about verdicts against the defendant in these cases.[153] In the word's of Stein, the evidence against the defendant in a case like the *Gatecrasher Case* 'is not weighty because most of the evidence that could verify or refute it is missing'[154] This can be even expressed in terms of second-order beliefs: the degree of our second-order belief, that additional inquiries would not decrease the probability of the hypothesis against the defendat below the standard of proof, is too low to find for the plaintiff on the existing evidence. Finally, David Kaye points out that in the *Gatecraseher* little is needed for additional evidence to change the probability of the hypothesis dramatically.[155]

Bus Case.

[153] (See 1999; 1509)

[154] (See Stein; 2005, 85)

[155] (See Kaye; 1979, 107) More fully, Kaye's view is that there is a conflict between objective and subjective probabilities. He begins by distinguishing between two probabilities by which we can analyze the naked statistical evidence in the Paradox: objective and subjective. Objective probabilities denote the chances or the odds that an event will occur. applied to the Gatecrasher, this probability indicates that each rodeo has a slightly greater than 50% likelihood of membership in the class of

tresspassers. Kaye also discusses problems with an unwavering commitment to just quantitative evidence. The hallmark of a trial, beit civil or criminal, is the evaluation of competing narratives. However, an over reliance quantitative evidence educes the trial's traditional narrative format to mere background statistics. Moreover, the accuracy decreases as the probative force of corresponding qualitative evidence is lost.

As will become clear, Kaye's view of subjective probabilities is more sanguine. The subjective probabilities combine quantitative evidence with a desire for corresponding qualitative support. So, a court could conceivably calculate subjective probabilities in two stages. It begins by calculating the objective probabilities and later discounts it in order to incentivize plaintiffs to do more than produce base rate evidence. It's easy to see the implications this might have on the Gatecrasher. A court calculating subjective probabilities might begin with the 50.1% objective probability that any given deffendant tresspasses, and then discount it in order to also require corresponding qualitative evidence. This strategy, were it to be used, would seem to surmount some of the unfairness concerns raised above.

Kaye's second response is more piercing. He argues that the 50.1% chance that any one of the thousand rodeo attendees is actually a gatercrasher is misleading. The probability, taken by itself, fails to take into consideration the possibility of exculpatory evidence that could be admitted but is withheld. As Kaye writes, "Hence, if it is even slightly more likely that the rodeo organizers would have been able to come forward with more evidence about how the defendant A came onto the premises without paying if he had actually done" (See Kaye; 103-107) so, the 50.1% chance that defendant A is a gatecrasher is an inflated mischaracterization. Thus, according

7.2.2 Injustice

7.2.2.1 Glanville Williams's view

"Weightiness considerations" are not the only factors that some scholars adduce as grounds for why liability should not be imposed in the *Gatecrasher*. Glanville Williams (1979), for example, believes that naked statistics is not sufficient to find for the plaintiff, because it is an injustice done to the randomly selected defendant.[1] Indeed, Williams argues that the fact-finder should never return a verdict for the plaintiff who does not tie the evidence to a specific defendant. This means that a plaintiff cannot simply choose a random individual out of a population of potential defendants. Instead, what the plaintiff needs in order to prevail is evidence tailored specifically to the defendant on the stand.

Williams further argues that the injustice created by the *Gatecrasher* could be eliminated through an apportionment system by which damages are awarded not solely based on the defendant's proportionate fault, but instead in proportion to that defendant's "fault-plus-probability." Such a system of damage apportionment would reflect the disconcerting link between the defendant and the naked statistical evidence used against him at trial. Additionally, it would reduce the award to a

---

to Kaye, a fairminded juror could hold two apparently conflicting beliefs. On the one hand, she could believe the naked statistical evidence against the Gatecrasher defendants. But, owing to external factors-such as withheld evidence of a particular defendant's innocence, she could also decrease the probability of that defendant's guilt to a figure below one-half.

plaintiff who relied solely on naked statistical evidence, incentivizing the production of additional evidence.[156]

7.2.2.2 Sir Richard Eggleston's view

Acting in accordance to putative principles of justice in the *Gatecrasher* does not always mean that the fact-finder must refrain from imposing liability. Sometimes, it may mean that the fact-finder *ought to* impose liability - at least that is Sir Richard Eggleston's position. Similar to Williams, Eggleston's analysis of the *Gatecrasher's Paradox* appeals to concepts like justice and fairness. But unlike Williams, Eggleston believes that the real injustice is done to the *plaintiff*. He writes: "Any injustice involved in giving judgment against the defendant in the rodeo case does not outweigh the injustice of refusing a remedy to the plaintiff who has the odds in his favor."[157] Eggleston acknowledges that there might be grounds to reduce or prevent the potential injustice for defendants by decreasing judgments. But, Eggleston focuses on the counter justice concern. According to Eggleston, by permitting attribution schemes in support of the defendants, we diminish the possibility of recovery of damages for the plaintiff. Thus, Eggleston finds no serious injustice done to the defendant in the *Gatecrasher's Paradox* and claims that naked statistical evidence is sufficient for a plaintiff to recover.

[156] (See Williams; 1979)

[157] (See Eggleston; 1980)

7.3 *LRA* and the Gatecrasher Paradox

7.3.1 LRA's solution

Our discussion so far considers several historical solutions to the *Gatecrasher Paradox*. Answers range from providing reasons for and against imposing liability on a random spectator, with the majority of the solutions siding with not imposing liability. The solution that *LRA* proposes also claims that liability should not be imposed on the defendant. To see this, consider what it would take, on *LRA*, for a randomly selected spectator to be held liable for gatecrashing. The theory requires that we start by determining whether a sufficiently large number of relevant alternatives to the claimant's narrative would be eliminated were defendant accused of gatecrashing. *LRA* provides two reasons for why it is very unlikely that the fact-finder would eliminate a sufficiently high number of relevant alternatives. The first seems obvious. Suppose a random spectator were accused of gatecrashing on the basis of bare statistics alone. The fact-finder would not be able to completely elim-inate the possibility that defendant is one of the attendees who *actually* paid for admission. This means then that the fact-finder would not find defendant liable.

Based on the features of *LRA*, there is another explanation it offers for why a pro-plaintiff ruling is unlikely. To see this, suppose again defendant is randomly selected from the population of spectators. He is subsequently tried. Finally suppose, at some point, the fact-finder is able to eliminate a fair number of alternatives, all of which support the conclusion that defendant did in fact gatecrash. According to *LRA*, we must then ask how resilient the fact-finder's belief about the defendant's apparent guilt would be if additional, albeit, conflicting, evidence were presented. If

we assume that the only evidence heretofore in support of the defendant gatecrashing is the base rate evidence presented in Section 7.1, then there seems to be many pieces of conflicting information that, if presented, would raise the probability of the fact-finder's belief changing. For example, suppose that there is a VIP section in the stadium designated only for those who paid cash in person and there is evidence that defendant viewed the sporting event in that section. Or, suppose defendant shows the transaction of the ticket purchase by pulling up his bank records. Any of these scenarios, if realized, would raise the possibility of the fact-finder's belief changing. Thus, if there is little (to no) resiliency in the fact-finder's belief that defendant gatecrashed, then this is a sign that a suciently high number of relevant alternatives have not been eliminated.

One further point needs to be mentioned here. While I maintain that Cohen's 1977 version of the Gatecrasher requires that the fact-finder refrain from imposing liability, that does not mean that all versions of the Gatecrasher must not impose liability on the defendant who is randomly selected. Given what I have argued above, there is in principle some cases of the Gatecrasher where it could be said that a suciently large number of relevant alternatives have been eliminated such that I would be committed to concluding that a pro-plaintiff liability verdict would be permissible. Consider, for example, a case in which 6,999,999.499 people paid for admission to some world event but that 7,000,000,000 are counted on the seats. I believe that randomly selecting a defendant could, in princple, produce a rational conviction that the defendant is at fault. Thus, not all Gatecrasher cases have the same result.

7.3.2 LRA and other views

Earlier in this chapter, I discuss several historical solutions to the *Gatecrasher Paradox*. I then argue that the solution *LRA* proposes aligns with what most scholars already think the outcome of the *Gatecrasher* should be. *Prima facie*, my last point may seem vapid, especially given the untold number of solutions that already reach the same conclusion that *LRA* comes to. What is unique, however, about *LRA* is that it can also illustrate the grounds for most of the proposed solutions in the literature. The next two Subsections illustrate how *LRA* can represent the *Keynesian weight* view as well as the injustice accounts.

7.3.2.1 LRA and Keynesian weight

Recall that according to the *Keynesian weight* view, cases that rely only on naked statistics suggest that the investigation was sloppy. Indeed, the cases claim that a serious investigation may not have been conducted. *LRA* can account for this assessment. We can imagine a scenario in which defendant is accused, tried, and convicted of gatecrashing based solely on base rate evidence. Defendant files for appeal and is given a hearing in appellate court. Suppose defendant argues that the decision of the lower court should be vacated on grounds that a serious investigation has not been conducted. Is not his defense a roundabout way of saying that not enough relevant alternatives have been ruled out? Indeed, suspicions of a shoddy investigation can often be attributed to hastily concluding that defendant is culpable before ruling out enough other persons of interest.

7.3.2.2 LRA and Injustice accounts

LRA also has the resources to represent each justice view discussed in this chapter. Recall Glanville Williams holds that an injustice is done to the defendant who is found liable on the basis of bare statistics. We can think of a scenario in the *Gatecrasher* in which defendant complains of an injustice after being found liable of gatecrashing. He learns that the only evidence the plaintiff adduces is naked statistics. He then rightly complains that it would be an injustice to find him liable, given that the evidence is not specific to him.

Upon reflection, defendant's contention is not unlike stating that a sufficient amount of relevant alternatives have not been ruled out. To claim that the evidence is not specific to the person is just to claim that the evidence does not uniquely pick out the person in question. But isn't this just another way of saying that the evidence does not eliminate all other possible alternatives? That is, it is to claim that the evidence does not differentiate the target item from other alternatives. If my analysis is correct, then it should be noted that *LRA* can easily capture the idea of eliminating relevant alternatives.

Finally, *LRA* can illustrate that there can be scenarios in the *Gatecrasher* where it would be morally permissible to impose liability on defendant. To see this, I return to the *Gatecrasher* case with a few modications (hereafter, modified *Gatecrasher*). Suppose this time that there are 1,000,000 attendees, instead of 1,000. Futhermore, all but two of the attendees do not pay for admission. Similar to the original *Gatecrasher*, let us assume that we just do not know which two attendees do pay. Suppose further that these bare statistics are all that is known about the case. If we assume

that it is never perimissible to impose liability on a randomly selected spectator, it follows then that liability would not be imposed on any of the spectators.

On the other hand, the likelihood that the randomly selected spectator *actually* gatecrashed seems really high and consequently cannot be overlooked. If ever there was a naked statistical evidence case that justifies imposing liability, this seems like a good candidate. Indeed, we can even see Eggleston's point. If the fact-finder chooses not to impose liability on any of the attendees because of qualms with relying on naked statistics, this can engender economic hardship on the plaintiff seeking to recover damages for the defendant's alleged, wrongful conduct. Consequently, plaintiff may rightly claim that the economic hardship is tantamount to experiencing an injustice, especially if the fact-finder refrains from imposing on grounds that some may seem unfounded.

What is important to note, however, is that *LRA* can capture the concerns raised by Eggleston. To see this, let us ask what it would take, on *LRA*, for a randomly selected spectator to be held liable for gatecrashing on the modified version. According to the theory, we ask the following question: Would a suciently high number of relevant alternatives to the plaintiff's narrative be eliminated were defendant accussed of gatecrashing? We might be tempted to think that *LRA*'s answer is *no*. However, this answer assumes that *LRA* requires *all* relevant alternatives to be eliminated in order to find for the plaintiff. However, such a standard is too exacting. *LRA*, instead, requires that a suciently high number of relevant alternatives be eliminated. Given that there is a 99,998 to 1,000,000 chance (i.e., a 99.9% probability) that defendant is liable, there are *prima facie* grounds for *LRA* to arm that enough relevant altternatives have been eliminated. Furthermore, the high likelihood also

suggests that not very many additional, yet conflicting, information, if presented, would change the fact-finder's belief that defendant is liable. Thus, *LRA* offers a straightforward explanation for why defendant in the modified *Gatecrasher ought to* be held liable.

References

[1] Aquinas, T. (1945). Basic Writings of St. Thomas Aquinas, trans. A.C. Pegis. New York: Random House.

[2] Bentham, Jeremy (1843) 'The Works of Jeremy Bentham, *vol. 6* (*The Rationale of Judicial Evidence*). William Tait, Edinburgh. Edited by John Bowring.

[3] Blome-Tilmann, Michael (2015) 'Sensitivity, Causality, and Statistical Evidence in Courts of Law' *Thought* 4(2): 102-112.

[4] Butler, J. (2005) *Forensic DNA typing.* Burlington, MA: Elsevier Academic Press.

[5] Carnap, R. (1962). *Logical foundations of probability* (2nd ed.) Chicago: University of Chicago Press

[6] Cauvin HE. Jury deadlocks in case that relied on DNA. Washington Post: Apr 15, 2006. Sect. B:10.

[7] Cheng, Edward K (2013) *Reconceptualizing the Burden of Proof*, 122 YALE L.J. 1254, 1269 - 71

[8] Cohen, L. Jonathan (1977) "THE PROBABLE AND THE PROVABLE. Oxford: Clarendon Press

[9] _ _ _ _ _ _ _ _ _ _ (1986) "The Role of Evidential Weight in Criminal Proof", *Boston University Law Review*, 66: 635-649

[10] Cohen, Stewart (1998). "Contextualist Solutions to Epistemological Problems: Scepticism, Gettier and the Lottery", *Australasian Journal of Philosophy*, 76: 289 - 306.

[11] Cohen, Stewart (1999). Contextualism, Skepticism, and the Structure of Reasons", *Philosophical Perspectives*, 13. J. Tomberlin: p. 57-90.

[12] Colorado v. New Mexico (1984), 467 U.S. 310.

[13] Colyvan, Mark et al., (2001) *Is It a Crime to Belong to a Reference Class?,* 9 J. Pol. Phil. 168, 171-172

[14] Convicted Murderer Seeks Retrial, Kalamazoo Gazette, May 10, 2006; Sect. Local News.

[15] Dahlman, Christian (2020) "Naked statistical evidence and incentives for lawful conduct". *Evidence & Proof*. Vol. 24(2) 162-179

[16] Dahlman, C; Wahlberg, L, and Sarwar F (2001) "Robust trust in expert testimony". *Humana Mente* 8: 17-37

[17] De Rose, Keith (1995). "Solving the Skeptical Problem", *Philosophical Review*, 104: 1-51.

[18] _____(1999) *Contextualism: An Explanation and Defense* in J. Greco and E. Sosa, ed., *The Blackwell Guide to Epistemology*, Blackwell Publishers.

[19] DeWitt, R. (2004). Worldviews: An Introduction to the History and Phi- losophy of Science. Blackwell Publishing Ltd.

[20] Di Bello, M (2013). Statistics and Probability in Criminal Trials: The Good, The Bad and the Ugly

[21] _ _ _ _ _ _ _ _ _ _ (2020). Proof Paradoxes and Normic Support: Socializing or Relativizing? Mind, Vol. 129. 516

DNA Exonerations Nationwide, INNOCENCE PROJECT, http://www.innocenceproject.org/Content/DNA_Exonerations_Nationwide.php

[23] Dretske, Fred (1970). "Epistemic Operators", *The Journal of Philosophy*, 67: 1007 - 1023

[24] _ _ _ _ _ _ _ _(1981). "The Pragmatic Dimension of Knowledge". *Philosophical Studies*, 40: 363 - 378

Dubse v. State, 1993 WL 381482 at *8 (Ala. Cr. App.)

[26] Enoch, David and Fisher, Talia (2015) Sense and 'Sensitivity': Epistemic and Instrumental Approaches to Statistical Evidence. Stanford Law Review, Vol. 67. No. 557

[27] Eggleston, R. (1980) *The Probability Debate*, 1980 CRIM. L. REV.

[28] Feinberg, S.E. (1986). *Gatecrashers, Blue Buses, and the Bayesian Representation of Legal Evidence,* 66 B.U.L. REV. 693, 693-98

[29] _ _ _ _ _ _ _ _ _ _ _ (1989). *The evolving role of statistical assessments as evidence in the courts.* New York: Springer-Verlag

Frye v United States 54 App. D.C. 46, 47, 293 F. 1013, 1014 (1923)

[30] Good, I. (1983). *Good thinking.* Minneola: Dover Publications.

[32] Gardiner, Georgi (2018), "Legal Burdens of Proof and Statistical Evidence" in *The Routledge Handbook of Applied Epistemology*, eds. David Coady and James Chase.

[33] ——————— (2019), "The Reasonable and the Relevant: Legal Standards of Proof" *Philosophy and Public Affairs* 47 (3):288-318

[34] Gettier, Edmund L. (1963) "Is Justified True Belief Knowledge *Analysis*. 23 (6): 121 - 123

[35] Gilmore, Carl W. *Challenging DNA in Paternity Cases: Finding Weaknesses in an Evidentiary Goliath*, 90 ILL. B.J. 472, 472

[36] Goodwing, Robert J. & Gurule, Jimmy (2002) Criminal and Scientific Evidence: Cases, Materials, Problems 333 (2d ed. 2002).

[37] Haack, Susan (1993). *Evidence and Inquiry: Towards Reconstruction in Epistemology*. Oxford: Blackwell.

[38] Harman, Gilbert (1986). Change in View. *Principles of Reasoning*. MIT

Jones v. Oklahoma, 554 P.2d 830, 835 (1976)

[40] Kagehiro, Dorothy K., Laufer, William S. (1992) Handbook of Psychology and Law. Springer-Verlag New York.

[41] Kaminsky v. Hertz Corp., 288 N.W.2d 426, 94 Mich. App. 356

[42] Kant, I. (1993). Grounding for the Metaphysics of Morals, trans. James W. Ellington. Hackett Publishing Company.

[43] Kasser, Jeff (2015). Two Conceptions of Weight of Evidence in Peirce's *Illustrations of the Logic of Science.* Erkenntnis 81, 629-648

[44] Kaye, D (1979). *The Laws of Probability and the Law of the Land*, 47 U. CHI. L. REV. 34, 36-39

[45] Kelly, T (2008) *Evidence: Fundamental Concepts and the Phenomenal Conception.* Philosophy Compass vol. 3, no.5, September 2008: 933-955

[46] Keynes, J. Maynard (1921) *A Treatise on Probability.* Macmillian

[47] Knight v. State, 435 S.E.2d 276, 278 (Ga. Ct. App. 1993)

[48] Laudan, Larry (2011) "The Rules of Trial, Political Morality and the Costs of Error: or, Is Proof Beyond a Reasonable Doubt Doing More Harm Than Good? In Leslie Green and Brian Leiter (eds) *Oxford Studies in Philosophy of Law* Oxford University Press, pp. 195-227.

[49] Lavoisier, A. (1862). Reflexions sur le Phlogistique in Oeuvres, vol. 2, 623-4 Paris: Imprimerie Imperiale.

[50] Lawlor, Krista (2014) Exploring the Stability of Belief: Resiliency and Temptation, Inquiry, 57:1, 1-27, DOI

[51] Lewis, David (1986) *Counterfactual Dependence and Time's Arrow* in 2 Philosophical Papers 32, 47

[52] _____ (1996) Elusive Knowledge, *Australasian Journal of Philosophy* 74: 549-67, in his *Papers in Metaphysics and Epistemology*, Cambridge: Cambridge University Press, 1999: 418-45.

[53] Lilquist, Erik *Recasting Reasonable Doubt: Decision Theory and the Virtues of Variability*, 36 U.C. Davis L. Rev. 85, 140 (2002)

[54] Lindsey, Samuel et al., (2003) "Communicating Statistical DNA evidence." *Jurimetrics* Vol. 43 No.2 pp. 147 - 163

[55] Lippke, Richard (2010) "Punishing the Guilty, Not Punishing the Innocent" *Journal of Moral Philosophy* 7(4): 462-88.

[56] Littlejohn, Clayton (2017) Truth, Knowledge, and Standard of Proof in Criminal Law. *Synthese*: 1-34.

[57] MacIntyre, A. (2006) The Tasks of Philosophy. Cambridge University Press.

[58] McCauliff, C.M.A. (1982) "Burdens of Proof: Degrees of Belief, Quanta of Evidence, or Constitutional Guarantees?" *Vanderbilt Law Review,* vol. 35(6): 1293 - 1336.

McCray v. City of New York 03 Civ. 9685 (DAB), 03 Civ. 9974 (DAB), 03 Civ. 10080 (DAB).

[60] Moss, Sarah (2021) "Knowledge and Legal Proof" forthcoming in *Oxford Studies in Epistemology* vol. 7: Oxford University Press

[61] Mulrine, Thomas (1997) "Fainthearted Conditionals' *Journal of Philosophy* 94(4): 187-211.

[62] Nesson, Charles (1979). Reasonable doubt and permissive inferences: the value of complexity. *Harvard Law Review,* 92(6), 1187-1225.

[63] ——————— (1985) *The Evidence or the Event? On Judicial Proof and the Acceptability of Verdicts*, 98 Harvard L. Rev. 1357, 1378.

[64] Newton, I. (1964). The Mathematical Principles of Natural Philosophy. New York: Citadel Press.

[65] https://www.nolo.com/legal-encyclopedia/legal-standards-proof.html

[66] Nozick, R. (1981). Philosophical Explanations Cambridge MA: Harvard University Press 172-185, 197-217

[67] Nunn, G. Alexander (2019) The Incompatibility of Due Process and Naked Statistical Evidence, 68 *Vanderbilt Law Review* 1407

[68] O'Donnell, R. (1989) *Keynes: Philosophy, economics and politics.* New York: St. Martin's Press.

[69] Pattenden, Rosemary (2009) "Pre-Verdict Judicial Fact-Finding in Criminal Trials with Juries". Oxford Journal of Legal Studies, Vol. 29, No.1 pp. 1-24

People v Jordan, 45 Cal.2d 697 (1955)

People v Malloy, 55 N.Y. 2d 296, 300, 303 (1982)

People v Soto, 35 Cal. Rptr. 2d 846, 847 (Cal. Ct. App. 1994)

[73] Peterson Joseph L. & Leggett, S Anna (2007) ; *The Evolution of Forensic Science: Progress amid the Pitfalls*, 36 STETSON L REV. 621, 630

[74] Popper, K. (1959). *The logic of scientific discovery.* New York: Harper

[75] Posner, Richard A. (1999) *An Economic Approach to the Law of Evidence*, 51 Stan. L. Rev. 1477, 1508 n.68

[76] Pundik, Amit (2008) Statistical evidence and individual litigants - a reconsideration of Wasserman's argument from autonomy. *International Journal of Evidence and Proof* 12: 303-324.

[77] Redmayne, Mike (2008). "Exploring the Proof Paradoxes." *Legal Theory*, vol. 14 (4): 281-309.

[78] Rich, Patricia (2020) "The logic of probabilistic knowledge." (2020) *Philosophical Studies* 177: 1703 - 1725

Riggs v. State, 809 N.E.2d 322 Ind. (2004))

Roberson v. State, 16 S. W.3d 156, 170 (Tex. Crim App 2000)

[80] Ross, Lewis (2019) Rehabilitating Statistical Evidence *Philosophy and Phenomenological Research.*

[82] Ross, Lee and Anderson, C. (1982) Shortcomings in the attribution process: On the origins and maintenance of erroneous social assessments. In Kahneman, Slovic & Tversky (eds), *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press. pp. 129-152

[83] Roth, Andrea (2010) *Safety in Numbers? Deciding When DNA Alone Is Enough to Convict,* (discussing the phenomenon of "pure cold hit" cases). 85 N.Y.U. L. REV. 1130, 1132-33

[84] Sanchirico, Chris W (2011) *Character Evidence and the Object of Trial*, 101 Colum. L Rev.

[85] Schaffer, Jonathan (2001) *Knowledge, relevant alternatives and missed clues.* ANALYSIS 61.3, July 2001, pp. 202-208.

[86] Schaffer, Jonathan (2005) *What Shifts? Thresholds, Standards, or Alternatives* in Gerhard Preyer and Georg Peter's (eds.), *Contextualism in Philosophy*: Knowledge, Meaning, and Truth, Oxford University Press, 2005

[87] Seeds, M. (1998). Horizons: Exploring the Universe. Wadsworth Publishing Company.

Sindell v. Abbott Laboratories 26 Cal. 3d 588 (1980)

[89] Skyrms, Brian. (1980) *Causal Necessity: A Pragmatic Investigation of the Necessity of Laws*. New Haven, CT: Yale University Press, 1980.

[90] Smartt, Timothy (2020) Reconsidering the rule of consideration: Probabilistic knowledge and legal proof. *Episteme, Early View Online First*.

Smith v. United States, 348 U.S. 147, 152 (1954)

[92] Smith, Martin (2016) *Between Probability and Certainty: What Justifies Belief*. Oxford University Press UK.

[93] _____ (2018) 'When Does Evidence Suffice for Conviction?' *Mind.* doi:10.1093/mind/fzx026

[94] Springfield v. State, 860 P. 2d 435, 449

State v. Briggs, 351 S. W.2d 892

State v. Hunter, 861 N.E.2d 898 (Ohio App. Ct. 2006).

State v. Pankow 144 Wis. 2d 23 (1988). 422 N. W.2d 913

State v. Toomes, 191 S. W. 3d 122, 131 & n.4 (Ct. Crim. App. Tenn. 2005)

[96] Stathis, P. (1999) Scientific Realism: How Science Tracks The Truth: Routledge.

[100] Stein, A. (2005) *Foundations of Evidence Law*. Oxford University Press

[101] Stine, G. (1976) Skepticism, Relevant Alternatives and Deductive Closure. *Philosophical Studies* **29** 249-261

[102] The Copernican Model: A Sun-Centered Solar System. http://csep10.phys.utk.edu/astr161/lect/retrograde/copernican.html

[103] The Innocence Project. Eyewitness Misidentification. (2014) Avail-able from: http://www.innocenceproject.org/understand/Eyewitness-Misidentification.php

[104] Thomson, Judith Jarvis (1986) *Liability and Individualized Evidence*, *in* Rights, Restitution, and Risk: Essays In Moral Theory 225, 230.

[105] Tribe, Laurence (1971) "Trial by Mathematics: Precision and Ritual in the Legal Process' *Harvard Law Review* 84(6): 1329 - 1393.

[106] *Turner v. United States,* 396 U.S. 398, 414-17 (1970)

[107] Ulian, J. (1995). On Peirce on induction: A response to Levi. In K. Ketner (Ed.), *Peirce and contemporary thought* (pp. 94-99). New York: Fordham University Press.

[108] *United States v. John Veysey*, No. 99 CR 381 (2003)

[109] *United States v. Jenkins*, 887 A.2d 1013 (D.C. 2005)

[110] *United States v. Rogers*, 617 F. Supp. 1024 (D. Colog 1985)

[111] *United States v. State*, 533 So. 2d 841, 850 (Fla. Dist. Ct. App. 1988); *see id.* at 843

[112] *Victor v. Nebraska*, supra 511 U.S. at 13 and 17-20

[113] Walen, Alec (2015) "Proof Beyond a Reasonable Doubt: A Balanced Retributive Account" *Louisiana Law Review* 76(2): 355-446.

[114] Walker, M. (2007). Moral Understanding: A Feminist Study in Ethics. Second ed. Oxford University Press

[115] Wasserman, David T. (1991) *The Morality of Statistical Proof and the Risk of Mistaken Liability*, Cardozo L. Rev. 935, 942-43.

[116] Wells, Gary (1992) "Naked Statistical Evidence of Liability: Is Subjective Probability Enough." *Journal of Personality and Social Psychology.* Vol. 62, No. 5, 739- 752

[117] Wheeler, G (2007). "A Review of the Lottery Paradox", in William Harper and Gregory Wheeler (eds.) *Probability and Inference: Essays in Honour of Henry E Kyburg, Jr.*, King's College Publications, pp. 1-31

[118] Whitman, James (2008). *The Origins of Reasonable Doubt: Theological Roots of the Criminal Trial.* Yalse University Press, New Haven.

[119] Whitmore, Damon (1993) *The Admissibility of DNA Evidence in Criminal Proceedings*, 39 WAYNE L. REV. 1411, 1411.

[120] Williams, Glanville (1979) *The Mathematics of Proof - I,* CRIM. L. REV. 297 - 301

[121] Williamson, T. (2000). *Knowledge and its limits.* Oxford: Oxford University Press

[122] Winship (1970), 397 U.S. 358, 364.

[123] Wright, Richard W. (1988) *Causation, Responsibility, Risk, Probability, Naked Statistics, and Proof: Pruning the Bramble Bush by Clarifying the Concepts,* 73 Iowa L. Rev. 1001, 1054.

[124] W3. *Writings of Charles S. Peirce: A chronological edition.* Bloomington: Indiana University Press.