The Foundations of Criminal Law Epistemology

Lewis Ross[1]
London School of Economics and Political Science

Abstract

Legal epistemology has been an area of great philosophical growth since the turn of the century. But recently, a number of philosophers have argued the entire project is misguided, claiming that it relies on an illicit transposition of the norms of individual epistemology to the legal arena. This paper uses these objections as a foil to consider the foundations of legal epistemology, particularly as it applies to the criminal law. The aim is to clarify the fundamental commitments of legal epistemology and suggest a way to vindicate it.

~~~

Introduction

Many philosophers have drawn conclusions about legal norms by appealing to the norms governing rational belief. The assumption that epistemology can tell us something important about the law—about the way the legal system *ought* to be—underpins the burgeoning project of legal epistemology, surely one of the notable philosophical trends of the last decade or so. By accepting some correspondence between legal and epistemic norms, we can understand legal verdicts as susceptible to evaluation by the very same normative categories used to assess beliefs. For instance, we might ask whether a given legal verdict counts as knowledge, surpasses some Lockean threshold for rational belief, or possesses various types of non-probabilistic justification. Our answers to these questions might then be used to vindicate or condemn aspects of legal procedure. In this sense, the promise of legal epistemology

---

[1] Comments welcome to L.ross2@lse.ac.uk

is to open philosophy of law to a fruitful exchange with the rich theoretical resources of individual epistemology.

But does the growing literature on legal epistemology rest on a basic philosophical howler? In a recent paper, David Enoch, Talia Fisher and Levi Spectre claim that it does.[2] They suggest that "this project is based on a mistake, roughly analogous to the mistake involved in thinking of studies of intelligence as relevant to the understanding of military intelligence."[3] Enoch, Spectre, and Fisher (henceforth 'ESF') demand an answer to the question: 'why we should ever sacrifice *accuracy* to ensure that legal verdicts possess some additional epistemic property?' They conclude that the prospects for answering this question are dim. In short, they charge legal epistemologists with exhibiting *epistemic fetishism*, by countenancing that we ought to prefer a less reliable legal system just so that legal verdicts instantiate a preferred epistemological talisman or 'Status'.

ESF are not alone in pressing doubts about the relevance of individual epistemology to legal philosophy. David Papineau (2021) calls preoccupation with legal *knowledge* a 'stone-age hangover' that interferes with accuracy. I (2021a) have challenged the purported symmetry between courts and individual believers. Giada Fratantonio (2021) suggests that legal epistemology is sorely in need of a 'value-turn'. And elsewhere, decision-theoretic approaches to legal adjudication are not sympathetic to the idea that we need to care about complicated epistemic properties, rather than simply relying on probability-weighted expected values.[4]

If we fail to answer this challenge, the entire project of legal epistemology is undermined. This paper uses the fetishism objection as a foil to consider the foundations of criminal law epistemology. My aim is to explain why legal epistemology—at least when applied to the criminal law—is on a solid footing. I will do so by reflecting on the role of the court in holding people to account and the conditions under which such judgements can be deemed legitimate by society. This approach, so I hope, will provide a justification for legal epistemology integrated into a broader

---

[2] See Backes (2020) for pessimistic consideration of the opposite claim—whether epistemology can learn from legal cases.

[3] Enoch, Spectre and Fisher (2021: 1).

[4] See Kaplan (1968) for a classic paper. For more recent work, see, for example, Lillquist (2002), Ribeiro (2019), or Di Bello (2019).

normative perspective on the role of criminal trials. However, the foundations I sketch will not write a blank cheque for legal epistemologists—as it will turn out, it may be that only projects that display a sensitivity to folk epistemology are especially relevant for the criminal law.

Legal Epistemology: The Guiding Assumption

It is undeniably tempting to draw a parallel between courts engaged in a criminal trial and individual agents who are aiming to settle some question. Both are occupied in a type of inquiry, consider evidence for and against propositions, engage in deliberation, and terminate inquiry by issuing a type of judgement. These broad similarities have long been recognised by those writing about the law. But what characterises recent work in legal epistemology is an attempt to use the precise machinery of analytic epistemology to theorise about legal practice.[5] So, for instance, it has been argued that certain legal verdicts ought to be *counterfactually safe* (e.g. see Pritchard 2018 or Pardo 2018), possess *normic justification* (see Smith 2018), be *sensitive* to the truth[6] (e.g. Günther 2021), eliminate *salient error-possibilities* (see Gardiner 2019), or amount to *knowledge* (e.g. see Blome-Tillmann 2017, Littlejohn 2020 or Moss 2018), among other theories. These theories go beyond offering general platitudes that legal verdicts must be reasoned or conscientious, instead arguing that legal systems should produce decisions that align with the (often rather exacting) epistemological requirements for various types of rational belief.

  Before moving on, we need to be a little more precise. The claim that these legal epistemologists are advancing is not—or at least, it should not be—that *all* legal verdicts should meet the standards for rational belief. One contestable example is the civil law, where the 'more likely than not' standard of proof seems to permit verdicts based on evidence that would not be strong enough to support an outright belief. For example, a positive civil verdict that is underpinned by a .51 credence is, on its face, correct, even if we assume that the standards for rational belief require that some

---

[5] See Thompson (1986) for an earlier, influential philosophical discussion.

[6] In an earlier landmark paper, Enoch, Spectre and Fisher (2012) themselves famously defend a sensitivity-based account. However, their argument is not that sensitivity matters *per se*, but that it only matters for instrumental reasons relating to creating incentives to follow the law.

Lockean threshold higher than .51 be surmounted.[7] This paper focuses on criminal law, not civil law. But it is even more clear in the criminal domain that it is a non-starter to suppose that *every* verdict must meet the standards for rational belief. This is because acquittals—'not guilty' verdicts—certainly do not require that the fact-finder rationally believes the person is not guilty. This issue will have real normative importance later, so it is worth explaining it with some care now.

Legal judgements are the product of a type of artificial reasoning that is rather different from the way in which regular beliefs are formed. Ordinarily, when an individual inquires into some question, they have the entire menu of cognitive options open to them: for example, they can believe *p,* they can believe ¬*p*, they can assign a .65 credence to *p*, they can assign a .25 credence to *p,* or indeed they can suspend judgement entirely. Conversely, a court can only find the accused guilty or not guilty: it cannot assign any type of credence, nor can it suspend judgement.[8]

Indeed, criminal judgements are not even the result of a free choice between 'guilty' and 'not guilty'. In a criminal trial, the fact-finder (i.e. jury, or the judge in a judge-only trial) is simply asked whether guilt has been proven 'beyond a reasonable doubt'.[9] If the answer is 'no', then a *not guilty verdict is returned as a default*, in lieu of guilt being proven beyond a reasonable doubt. This technical point is important: it means that a not guilty verdict unequivocally does *not* entail that the fact-finder believes that the accused is innocent. For example, the evidence might weakly suggest guilt, but numerous reasonable doubts remain in the eyes of the jury. Here, a not guilty verdict

---

[7] Of course, many in the legal epistemology literature defend additional non-probabilistic conditions on appropriate civil verdicts. However, meeting these non-probabilistic requirements does not mean that a civil verdict must be believed outright in order to be appropriate. To take one example, one can have probabilistic knowledge of the likelihood of a proposition in the sense of Moss (2018) without believing that proposition outright.

[8] In Scotland, there is a third verdict: *not proven*. From the perspective of the law, it is functionally equivalent to 'not guilty', although this tripartite verdict structure is more informationally rich. What I say here still applies to such a system, namely that the options open to the court are constrained, and rational belief is not required in the content of every verdict.

[9] The 'beyond reasonable doubt' formulation is the one found in common law countries, but the same general idea applies in civil law traditions which formulate the criminal standard of proof in different terms.

would (rightly) be returned, even though each individual juror might think it more likely than not that the accused is guilty.

The foregoing is simply a descriptive truism about how trials operate, but there are good reasons for *not* requiring juries to believe that the accused is innocent before acquitting. Administering justice would be impossible if we could only deliver a verdict when the evidence decisively supported either guilt or innocence: sometimes evidence is equivocal. What do we do in these equivocal cases? Many have supposed that convicting the innocent is worse than acquitting the guilty. So, putting these points together, the law's solution is to impose a demanding burden of proof on the prosecution to rebut the presumption of innocence that we apply to the accused. When this burden of proof is not discharged, and the evidence remains equivocal between guilt and innocence, a not guilty verdict is returned. In light of these considerations, it cannot be a normative expectation that 'not guilty' verdicts meet the standards required for rational belief that the accused is not guilty. For example, it would be entirely appropriate to return a not guilty verdict if each individual juror rationally suspended judgement on the guilt of the accused.

Given this, a natural way to view the legal epistemology project is as arguing for a correspondence between the norms for 'guilty' verdicts and rational beliefs. I want to introduce a name for this position, which I will dub *legal doxasticism*. Stated more precisely, the idea is as follows:

> **Legal Doxasticism:** Guilty verdicts are appropriate only if full belief in guilt would possess a specific rationality-conferring property, given admissible evidence.[10]

The idea that guilty verdicts should correspond with the norms of epistemology is no obvious conceptual truism, even regarding the necessity rather than sufficiency version of the doxasticist norm stated above. There are clear differences between legal verdicts and beliefs, which leave us in need of a further normative story for why we should treat the two as if they are governed by similar norms. Most obviously, guilty

---

[10] A simpler version of the norm would be: 'only if a full belief in guilt *would be rational*, …'. This would capture most views in the literature, where proponents also disagree about the norm of belief. But, we should be open to the idea that there are a plurality of ways in which a belief can be rational and that only some of these ways might matter for legal epistemology.

verdicts are made with explicit reference to a standard of proof—'beyond reasonable doubt'—and standards of proof are something that can be deliberately amended. (Indeed, we've just explained that civil law seems to use a standard of proof that falls below the standard for rational outright belief). By contrast, the standards for rational belief are often taken to be given a priori or otherwise immune from conscious at-a-moment revision. So, why should we set the standard of proof so that is at least as demanding so as to require that the judge or jury rationally believes that the accused is guilty before convicting them? This is the normative question that I want to answer, and, in so doing, clarify and challenge ESF's argument that it is a mistake to suppose that the law should care about the standards of individual epistemology.

The Fetish Objection

ESF charge legal epistemologists with fetishism. They argue that it is misguided to subject legal discourse to the idiosyncratic epistemic proclivities of philosophers, when the truth is that we should never countenance an *accuracy-sacrifice* just to secure the presence of some additional epistemic status (such as knowledge, safety, normic support, etc.). These epistemic statuses, so the objection goes, have no intrinsic legal value when we are careful to separate them out from any instrumental connection with accuracy. (For instance, perhaps aiming for knowledge helps make legal fact-finders more accurate, even if knowledge is not intrinsically valuable). If you *really* hold that some outcome is intrinsically valuable, ESF write, "then you must be willing to pay, at least sometimes, at least some price in the hard currency of […] other goals."[11]

  The relevant goals they see as occupying the other side of the scales are the accuracy-goals of: (a) convicting only the guilty, and (b) acquitting only the innocent. Concerning the first of these accuracy-goals, only convicting the guilty, ESF write (using 'knowing' as a placeholder for any favoured epistemic status):

> How many more false convictions are you willing to allow this system to generate,
> just in order to make sure that the system does better in terms of *knowing* its
> findings than some alternative system (that doesn't care about knowledge, but is a
> bit more accurate, and renders a few less false convictions)? [Enoch, Spectre and
> Fisher 2021: 7]

---

[11] Enoch, Spectre and Fisher (2021: 5).

On its face, this challenge does not strike me as forceful. It is hard to envisage cases where demanding an *additional* epistemic safeguard before convicting will generate more false convictions than failing to demand this safeguard. The general idea of imposing an additional epistemic requirement on criminal conviction is to remove rather than add cases from the 'okay to convict' pile. As far as I understand those who argue that we should impose epistemic requirements on criminal conviction, they do not mean to say that we should replace the existing standards of criminal proof with weaker standards.[12] Rather, imposing epistemic requirements on conviction would be a way to tighten our standards, leading to less mistaken convictions rather than allowing for more.

---

[12] Going into the weeds a little further, this objection would only have force against a very specific type of view, one I am not sure anyone has defended. To understand the type of view the objection could apply to, note the following two ways to think about standards of proof:

1. *Probabilistic Support*: Convict only if the likelihood of guilt is $> n$.

2. *Non-probabilistic support*: Convict only if the verdict would possess the following non-probabilistic epistemic property...
   a) It would be safe, or
   b) It would be normically supported, or
   c) It would be sensitive, or

   ...etc.

As phrased, these conditions are necessary for appropriate conviction while remaining silent whether they are sufficient.

The only way that privileging some epistemic property could ever be expected to lead to *more* false convictions is if we *replace* the idea that surmounting any probabilistic threshold is necessary for conviction with the idea that it is *sufficient* that the verdict possess a certain type of non-probabilistic justification, where this type of non-probabilistic justification does not entail a likelihood of correctness greater than n. It is true that there are types of non-probabilistic justification that can apply to a belief while failing to guarantee that it meets a certain probabilistic threshold. Normalcy views might be one such example, where something can be normal in Smith's (2018) sense of 'calling for special explanation' when it fails to occur, even if the thing fails to occur fairly often. (For example, think about a dilapidated old car failing to run smoothly on various days). However, as I understand proponents of views with this structure, their discussions are aimed at defending necessity rather than sufficiency norms for criminal conviction.

The second objection concerning the problem of *only acquitting the innocent* is much more challenging. Turning to false acquittals, ESF put their argument against the intrinsic importance of these additional epistemic properties—which they dub 'The Status'—in stark terms:

> How many more people are you willing to have assaulted, or murdered, or raped under your designed system, just in order to secure [The Status] for the findings of your criminal justice system? [Enoch, Spectre and Fisher 2021: 5]

This challenge *is* forceful, because there will plausibly be situations in which imposing additional epistemic requirements on criminal conviction will lead to guilty people being acquitted who otherwise would have been convicted, absent these requirements. These situations are most clearly brought out by proof-paradoxical scenarios which— by and large—motivated much of the legal epistemology literature over the last ten years. Proof paradox cases are those in which we have overwhelming probabilistic evidence, *solely* in the form of some inculpatory statistic, tying some (group of) person(s) to a crime.[13] Here's an example of such a case:

> **Prisoners**. One hundred prisoners are exercising in the prison yard. Ninety-nine of them suddenly join in a planned attack on a prison guard; the hundredth prisoner plays no part. There is no evidence available to show who joined in and who did not. [Adapted from Redmayne 2008]

On a purely probabilistic conception of criminal proof, such evidence should suffice for convicting an individual prisoner, given the strong degree of probabilistic support for guilt. But many have a strong intuitive resistance to this idea.[14] According to *legal doxasticism*, conviction will be blocked in these cases, because beliefs concerning propositions with only statistical support will not possess the relevant epistemic status. (Just like beliefs in lottery propositions are not safe, sensitive, normally

---

[13] Many of the examples discussed in the literature are fairly artificial (see Pardo 2019 or Ross 2020 for a summary of work). But DNA evidence seems to be an example of a purely statistical type of evidence that is commonly found in criminal trials and which some think should be enough, by itself, to convict somebody of a crime (see Roth 2010 for discussion).

[14] However, this intuitive reluctance can arguably be disrupted by considering cases with multiple sources of statistical evidence, see (Ross 2021c). Section 5 of that paper surveys some empirical results relating to the proof paradox.

justified, known, and so forth). This has been taken by many to be a powerful argument in favour of epistemic requirements on criminal conviction.[15]

However, unlike individual agents considering lottery tickets, courts cannot hedge their bets and decide that they will wait and see what will happen—the court has an obligation to either convict or acquit. A doxasticist norm requires acquittal in these statistics-only cases, despite the overwhelming likelihood of guilt. So, ESF are quite right to say that privileging certain epistemic statuses will plausibly lead to more mistaken acquittals in some cases compared to a system that convicts only based on probability. And, if we mistakenly acquit more people, then some of these people will foreseeably go on to commit further crimes such as assault, murder, and rape, just so we can secure our preferred epistemic status for verdicts. Why on earth would we countenance this?

Responding to this challenge is the challenge of providing a justification for the project of criminal law epistemology.

Accuracy and Utility

Readers well-versed in their revisionary legal philosophy might notice that the rhetoric used by ESF—appealing to the risks of recidivist criminals murdering and raping—is very similar to that used by Larry Laudan, who famously argues against the demanding 'beyond reasonable doubt' standard of criminal proof. Laudan instead suggests that we ought to use considerably lower standards of proof when deciding whether to convict someone of a crime. Laudan, drawing on his interpretation of statistics[16] pertaining to violent criminal recidivism, writes that:

> [One is much] more likely to be the victim of a violent crime than to be falsely convicted of one. Unless we hold that it is preferable to be raped fifty times rather than falsely convicted once of rape, or that thirty murders [...] are preferable to one [person] being falsely convicted of murder, it should already be clear that the

---

[15] Many have also argued for similar epistemic requirements on civil findings. I do not discuss these arguments here, as I hold that the requirements of criminal and civil proof require different philosophical treatment.

[16] Note that Laudan's use of statistics has been subject to a comprehensive critique in Gardiner (2017).

marginal utility of taking steps to further reduce false convictions is much less than the marginal utility of reducing crime. [Laudan 2011: 199]

Observe, however, that Laudan puts his point not in terms of *accuracy* (as do ESF) but rather in terms of *utility*. This is surely sensible, if one is going to rely on the types of considerations that both Laudan and ESF appeal to. The reason that accuracy matters, on their accounts, is because inaccuracy leads to serious harms occurring as a result.

Indeed, once we appreciate this point, it quickly becomes clear, by talking so consistently about accuracy, that ESF are themselves indulging a mild fetish. The *intrinsic* disvalue of legal inaccuracy is not particularly weighty; any disvalue is dwarfed by the consequences of the false convictions and false acquittals that flow from inaccuracy.[17] (To see this, suppose that criminal cases were heard in secret, on a small island, in a beachside courtroom filled with palm fronds and rattan furniture, and where verdicts were written down, shown disinterestedly to the parties involved, and then used by the judges only to sop up spilled Mai Tais, with no further consequences.) Of course, the fact that it is not inaccuracy *per se* that matters is implicitly recognised by ESF, given their focus on the downstream harms—the murdering and assaulting—that flows from mistakenly releasing the guilty.

However, a little reflection reveals that promoting utility will not always march in lockstep with the global accuracy of the legal system. Let us say that the global accuracy of a legal system is the proportion of cases decided correctly/decided wrongly. Scenarios where global accuracy and promoting utility diverge are easy to conceive of. For instance, suppose that it were possible to make some procedural change which would have the effect of imprisoning one Very Dangerous criminal but would also result in two innocent people being fined twenty dollars for some petty infraction they did not commit.[18] This would promote utility, since the disutility of two mistaken $20

---

[17] The importance of non-accuracy related values is of course recognised at the end of Enoch, Spectre and Fisher's classic (2012) paper, where they defend a sensitivity condition on criminal conviction by appealing to the positive effects it has on incentivising law-abiding behaviour and deterring crime.

[18] This is stylised and abstract for expository purposes, but legal systems often face choices about whether to remove procedural protections that will affect certain crime-types more than others. For example, certain rules of evidence that require multiple sources of evidence before convicting might be harder to satisfy in sexual crimes compared to other types of crime. The Scottish legal debate about 'corroboration' is instructive in this regard.

fines is surely outweighed by the murders and rapes that the Very Dangerous person would go on to commit if mistakenly released.

The divergence between accuracy and utility-promotion could, in principle, be substantial. The disutility of convicting the innocent has some value, determined predominantly by the severity of the punishment to which they are subjected. Similarly, the disutility of acquitting the guilty has some *expected* value, determined by, inter alia, their risk of reoffending and the harmfulness of the recidivist crime they might commit. These are entirely contingent and changeable variables. So, the relationship between accuracy and utility is a contingent matter that turns on a variety of factors, including just how harshly we punish those we mistakenly convict, and just how dangerous are those who we risk mistakenly acquitting. Laudan's argument is alive to the empirical dimensions of this relationship when he argues that, in present actual societies, we might be better off convicting more people than we do at present. Other legal scholars have argued similarly, for example Epps (2015) argues that convicting more readily will promote utility by reducing perverse incentives that criminal have to plead not guilty—and subject society and victims to costly and prolonged trials—when accused. And elsewhere, some have argued that utility might be better promoted by having different standards of proof for different crimes, where we accept a lesser degree of accuracy in some contexts, in order that those who pose the most danger to society are incapacitated.[19]

So, where have we got to so far? If we accept the force of utility-promoting considerations in setting the standards of criminal proof, then ESF's focus on accuracy itself appears to be a fetish, or, at least, a red herring. The relationship between the global accuracy of a legal system and the extent to which it promotes utility is not straightforward, and the two values will predictably diverge in some cases.

What does this mean for the relationship between guilty verdicts and rational belief? ESF motivated their initial criticism by taking aim at the various *non-probabilistic requirements* that epistemologists have thought to be conditions on rational belief, requirements such as normalcy or safety. But their argument turns out to equally undermine the idea that criminal verdicts must necessarily be justified above whatever *probabilistic threshold* is required for a belief to be rational. For, it is entirely possible

---

[19] For example, see arguments in Ribeiro (2019).

for the contingent empirical facts to be such that utility would be best promoted by a standard of criminal proof that was lower than the Lockean threshold on rational belief. E.g., in a society with especially dangerous and recidivism-prone criminals, perhaps a standard of proof as low as .5 might best promote utility, despite this not being a sufficiently high credence to undergird a rational belief. Working out the best standard of criminal proof would, under that picture, simply be a matter for decision-theory: the answer would not be constrained by any doxastic norms, probabilistic or non-probabilistic. Indeed, that we have ended up here is no real surprise, given the tenor of ESF's core argument that epistemic constraints on legal norms are fetishes.

What I have done so far is to present the upshots of really taking seriously the idea that the norms of criminal convictions should be determined by the factors that ESF enjoin us to consider. This has not been an argument against the view as such, although I suspect many readers might already feel extreme discomfort with the idea of convicting someone of a crime based on a middling credence. Now, I want to explain why I think we should reject this view, and instead hold firm to the doxasticist idea that the standards of criminal conviction should be at least as demanding as those of rational belief.

Legal Epistemology and Moral Accountability

Why should we not convict someone of a crime unless we fully believe that they are guilty? Can we answer this question without appealing to the comparative disutility of different types of legal mistakes (an approach that will not necessarily vindicate legal doxasticism)? I will attempt to do so, but firstly I want to note one style of response that I do not think will provide the type of answer we are looking for.

An obvious way to take issue with the idea of convicting someone of a crime on a middling credence would be to embark on the project, familiar to any student of normative theory, of teasing out why it is generally problematic to jettison individual rights to promote utility. For example, we might focus on the idea that those accused of crimes are entitled to certain procedural safeguards as a ward against false conviction. This approach may be instructive. But it would not provide a clear route to showing that the right standards of criminal proof have any essential connection to the standards of rational *belief*. (For, one might suppose that we can protect rights simply by using an extremely high probabilistic threshold for conviction, without invoking the

standards of rational belief as an explanatory wheel). In other words, such an argument would not tell us anything about the foundations of legal epistemology, because it would not support legal doxasticism. We therefore must look elsewhere.

<div align="center">The Legal-Interpersonal Bridge</div>

In recent literature on the problem of basing a criminal conviction on merely statistical evidence—a literature ESF are reacting against—various authors refer approvingly to Lara Buchak's point that *blame is a belief-normed practice*.[20] In Buchak's view, one of the distinguishing hallmarks of belief versus credence is that only the former is an appropriate basis on which to hold someone morally responsible (i.e. view them as the legitimate object of various Strawsonian reactive attitudes). It is, seemingly, a truism about interpersonal attributions of responsibility that we only blame people for wrongdoing conditional on believing they have done wrong. (For example, consider how the statement: 'I blame you for *x*, but don't believe that you are responsible for *x*' seems to misfire). We do not proportion our blame to our credence: for example, a .6 confidence that someone has erred does not lead us to blame them to .6 of the degree to which we would blame them if we had credence 1 in their wrongdoing.

However, it has not yet been made explicit whether Buchak's idea tells us anything about the foundations of criminal law epistemology. For, it is not clear why the legal system ought to have norms that mirror the belief-normed nature of interpersonal responsibility practices.[21] To vindicate legal doxasticism, the connection between the interpersonal and the law is the bridge that must be defended.

> **The Legal-Interpersonal Bridge:** The legal system should care about the norms of interpersonal responsibility-attributions.

As ESF might pose the question, what do we lose if we sever the connection between interpersonal and legal norms? I think the answer, in rough terms, is the following: a criminal justice system that severs the connection with our interpersonal practices can no longer be viewed as punishing or holding people morally accountable on our behalf. Rather, it can only be viewed as a system of risk-management. I will explain what I

---

[20] See Buchak (2014).

[21] To my mind Littlejohn (2020), in which he aims to defend a focus on legal knowledge, comes closest to doing so.

mean, in order that we may better consider whether caring about moral responsibility in this way is a type of fetish.

Societies use the threat of force to institute various types of regime that promote general utility. To take a topical example, consider compulsory quarantine. Governments typically possess the power to compel those suspected of having dangerous contagions to sequester themselves until such time as they are no longer a risk to others. Clearly, when it comes to safeguarding society from dangerous contagions, a 'doxasticist' norm would not be an ultimately plausible condition on state-enforced quarantine. For instance, if our best tests could only establish to a .5 probability that a given individual has a virulent and potentially fatal disease, then this ought to be no bar to enforcing their quarantine. There is no obvious analogous 'proof paradox' dilemma in the context of enforcing quarantine for dangerous contagions.

Clearly, it would be an obvious error to say that we are *holding to account* or *punishing* the person whom we are forcing to quarantine. It would not be punishment, even if they are dragged away by the same agents of the state who take convicted persons to be imprisoned, and even if they undergo the same treatment as some penal regimes impose on convicted offenders, namely being confined to small cell. This is a general point, not confined to medical cases. For instance, pre-emptive incarceration of certain groups—e.g. internment during wartime—might even be administered by the same authorities as the penal system, yet such ill treatment should not be regarded as a way of holding someone morally accountable for their actions.

In contrast to these cases, punishment and the type of accountability found in the criminal law are responses for *culpably breaking* a rule or norm. Most crimes come with a *mens rea* requirement, which can roughly be understood as a culpable mental state (e.g. intending a homicide, or being reckless as to whether one would eventuate). Beyond the *mens rea* requirement, criminal law is preoccupied with whether the accused has any type of excuse or 'special defence' that diminishes or entirely undercuts their culpability.[22] For instance, those acting under compulsion or from insanity are excluded from criminal conviction, even if causally responsible for the

---

[22] Even those jurisprudentially awkward offences that are solely act-based—i.e. strict liability offences— generally allow for excuses, demonstrating that even regarding these offences the court is still aiming to establish moral responsibility.

physical component of a crime (e.g. causing a death). In a nutshell, then, punishing and holding to account involves attributing wrongdoing to someone as a morally responsible agent.[23] Deciding whether someone is morally culpable for wrongdoing is the bread-and-butter of criminal justice.

This form of treatment, one that is concerned with culpability, is entirely different from treating someone *simply as a vector of risk*, which is what we do in mandatory quarantine, pre-emptive incarceration, and restraint of those without the mental capacity to control their behaviour. When we treat someone as a vector of risk, we take a purely external perspective on their behaviour, treating them like a mechanism to be predicted and managed. We are free to justify our behaviour based only on the amount of risk they pose with no regard for moral responsibility. (It would be absurd, for example, to demand those in mandatory quarantine to offer *apologies*). Treating someone merely as a vector of risk is an entirely different mode of justification than that used by the criminal law. The criminal law engages those it finds guilty with *moral address*, informing them that they are being held to account by society for violating a standard that they were expected to conform to.[24] It is a prerequisite of this type of moral address that it is follows an investigation into whether the agent is morally responsible for the crime in question.

What is so instructive about Buchak's point concerning the belief-normed nature of moral responsibility attributions is that they are all-or-nothing. Treating someone as culpable for wrongdoing dovetails with what we believe about what they have done, rather than being graded in proportion to our confidence that they pose a threat. Putting this together with the point that punishment and holding to account are by their nature practices that involve attributing moral culpability, we see why it is no fetish for a criminal justice system to care about belief. Working out who is morally culpable for norm-violation is part of the distinctive role of criminal adjudication. This focus on moral culpability is what sets criminal justice apart from other state-backed

---

[23] Of course we can punish the innocent, perhaps intentionally. But there must be a background norm of aiming to punish the guilty in order for these instances to count as a species of punishment. (Just as we can tell lies when making assertions: this only makes sense against a background of a general expectation of truth-telling).

[24] The terminology of 'moral address' is due to Gary Watson.

regimes of adverse treatment which invoke only the notions of risk and utility-maximisation.

Judging on Our Behalf

I want to consider an important objection. The objection runs as follows:

> **Objection**. 'Courts are not bound by the same psychological constraints as individuals! Criminal courts are free to attribute responsibility and punish on the basis of whatever norms are adopted by the institutional framework: these can include norms which attribute moral responsibility without individual belief entering the picture. Just as civil verdicts do not require full belief, nor need criminal verdicts. The connection with our interpersonal norms should be severed if such a change would maximise utility.'

There is something correct in this objection. It is true that courts are not bound by the same constraints as individual agents when adopting attitudes or promulgating judgements. Indeed, as has been demonstrated in work on the 'discursive dilemma' and associated impossibility theorems, corporate agents can enact rules that allow them to issue judgements and form attitudes *entirely* independent from what their individual members believe.[25] Corporate bodies—for example courts, panels, government agencies, academic committees—can make judgements involving concepts which are belief-normed in individual agents, without being bound by the same psychological constraints as individual agents using these concepts. For example, suppose that finding a violin recital 'beautiful' is a belief-normed practice in individuals. You can only judge a violin recital to be beautiful if you believe that it is.[26] Nonetheless, a *School Awards Panel* might have a rule which means that they are constitutionally bound to hand out one 'beautiful recital' award per competition. The rules might dictate that the Panel should give the award to a recital that no member of the Panel individually found beautiful, but which nonetheless came closer than any of its rivals to being beautiful. The Panel would then make an award that confers a judgement involving a belief-normed practice in individuals (beauty-ascription), even if the individual panel members acknowledge that the evidence doesn't support a full

---

[25] For example, see List and Pettit (2002), List (2006), Pettit (2007).

[26] This seems fairly plausible to me, but I won't argue for it. If you are not convinced, substitute any other concept that you think is belief-normed in the individual and the argument will go through.

belief that the recital in question was beautiful. The same is true with courts—they can have rules that generate responsibility-attributing guilty judgements, without the evidence supporting (full) rational belief of the accused's responsibility in an individual agent, just so long as whatever rules that they have chosen to license guilty verdicts are fulfilled.

However, while the technical credentials of this objection are good, reflecting on it reveals why legal doxasticism is compelling. The nub of the point is as follows. Courts judge on behalf of a society. It is important that criminal guilty verdicts are in harmony with the would-be judgements of those the court is supposed to represent, namely the member of the society in which the court operates. In other words, it is important that criminal guilty verdicts are an effective *social signal* that, if you had considered the admissible evidence for yourself, you would agree that the accused is guilty. Criminal judgements—judgements of moral responsibility that license punishment—*would ring hollow if members of society could not be in a position to share them.* Even if courts could use their institutional levers to assign moral responsibility without the evidence supporting full belief in guilt, those whom the criminal court represents do not. It is, I suggest, a prerequisite for legitimacy that a criminal justice system only holds to account those whom a reasonable member of society could themselves deem culpable, if they were to consider the issue for themselves. A court which fails to uphold the doxasticist principle fails to attribute moral responsibility in conditions that would be upheld by a reasonable member of the community.[27] This is why it is important for criminal verdicts to satisfy the doxasticist principle. Conforming to legal doxasticism is required for the judgements of a criminal court to represent society in licensing blame-expressing punishment, and not merely a system of risk-management. Without satisfying legal doxasticism, criminal justice would no longer be regarded as a distinctive type of inquiry into culpability, but rather just as a coercive machinery for dealing with vectors of risk.

Of course, one might argue that caring about moral responsibility itself is a fetish. Perhaps we could axe the idea that there is any importance in judging people morally responsible before punishing them? I am inclined to think that caring about moral

---

[27] There is a corollary of this point, concerning the wrongdoer. While I lack space to develop this idea here, I believe that people are less likely to reconcile themselves to their treatment if they have been treated in a way that diverges from interpersonal norms.

responsibility in criminal justice is not a fetish, but rather a foundational idea. All normative theorising must accept some axioms or at least defeasible principles from which to proceed. The idea that 'we should only punish those we judge morally responsible for wrongdoing' seems to me like a strong candidate for such a principle. That we should presuppose such a principle is not something that can be conclusively demonstrated, but I think that it is much harder to call moral responsibility a fetish than the various epistemic concepts that have preoccupied legal epistemologists. If we to stop seeing the criminal justice system as an effective arbiter of moral responsibility, our relationship with it would thoroughly change. And if moral responsibility *is* important for criminal justice, then the Legal-Interpersonal Bridge can be defended: some correspondence between interpersonal and legal norms is required for the criminal justice system to legitimately adjudicate moral responsibility on our behalf.

Interpersonal Responsibility without Belief?

I have defended the project of legal epistemology by appealing to facts about when individuals attribute responsibility—namely only on the basis of a full belief—and suggesting that these facts constrain when guilty verdicts issued by a criminal court can be viewed as legitimate. This Legal-Interpersonal Bridge provides a foundation for the project of criminal law epistemology, or so I have argued.

But what if this foundation is rotten? What if our *interpersonal* blaming practices are themselves open to criticism? In brief yet suggestive remarks, David Papineau suggests exactly this. Papineau argues[28]:

> *We should reform our blaming practices as well as our legal ones:* we should abandon the shibboleth of not blaming on statistical grounds; this is driven by nothing but a mistaken concern with the outmoded category of knowledge; it only stands in the way of our distributing blame where it is most deserved. [Papineau 2021: 5325]

Papineau then continues:

---

[28] Papineau is technically only talking about knowledge, but his point readily generalises to any non-probabilistic condition on rational belief and his paper can be read as an attack on the entire legal epistemology project.

> It is not at all obvious to me that our lives would be much altered if we ceased to
> insist that personal reactions demand knowledge [...] After all, it is not as if this
> change would somehow undermine all the reactive attitudes. We could continue
> to blame and praise and so on, even if we came to do so on slightly different
> grounds. [Papineau 2021: 5326]

I don't know if I share Papineau's conviction that blaming without belief would only amount to an insignificant alteration to our lives. Certainly, if the idea is to maximise 'deserved blame' as suggested by the first quotation, then we will face difficulties stemming from the *gradable* nature of desert. One person might be extremely blameworthy (say for a terrible misdeed) and another only a little blameworthy (for minor wrongdoing). How should these differences in desert affect the level of confidence needed to attribute responsibility? And, if maximising the probability of allocating deserved blame required fine calculation or the application of different standards to different people, would this approach to blaming strike us acceptable?

Papineau doesn't answer these questions, answers for which are needed to substantiate and test his suggestion. But these questions are not ones we can answer here, as they would require a wide-ranging discussion of (i) interpersonal relationships and the conditions in which they flourish, and (ii) the comparative badness of different interpersonal blaming mistakes (false positives and false negatives in a variety of different contexts). However, we can say something about the jurisprudential ramifications of what Papineau is suggesting.

There are different ways one might develop Papineau's theme. The first way to interpret Papineau's remarks is to suppose that we should *first revise our interpersonal responsibility practices and then reform the legal system to follow suit*. Where would this proposal leave the project of legal epistemology? In the short run, legal epistemology would remain untouched: any change would await the (presumably lengthy) project of revising our interpersonal norms. In the longer term, the proposal would entail that we reject legal doxasticism, since the basic suggestion is that rational belief should not be required for blame, legal or interpersonal. However, this proposal would still retain some type of legal epistemology, since we would be retaining the idea that guilty verdicts should be founded on standards *at least as strong as those used by individuals*. So, this view is no threat to criminal law epistemology *tout court*, as the Legal-Interpersonal Bridge would be retained. But it would change the character of

legal epistemology and require replacing the doxasticist norm with whatever other notion ought to govern our individual responsibility-attributions.

A different way to interpret Papineau's suggestion has more immediate jurisprudential bite. This second interpretation is that courts should act as a sort of revolutionary vanguard, that we should *first revise our legal practices in the hope that our interpersonal practices will then follow suit.* This view *is* a real threat to criminal law epistemology because it endorses the idea that guilty verdicts can be appropriate even when individuals in society would not tend to share them. But should courts ever take on such a reformist role? This is a difficult question. Certainly, I do not think that courts need to wait until a guilty verdict would command *universal* assent from the society it represents. There will always be morally benighted persons who are slow to discard outmoded perspectives—such as misogynistic views relating to sexual criminality—and justice should not wait for these people to be convinced. However, I *do* think that moral reforms should go through the mill of social debate and have a modicum of acceptance before it is reasonable for a court to begin to enforce them. Changing our interpersonal blaming practices from a belief-normed to a probability-centric approach strikes me as a radical shift (indeed, one might wonder how psychologically possible such a shift is). It is not the role of a professionalised judiciary to promote such radical shifts without the ideas underpinning these shifts first being scrutinised at some length. I do not think that this has yet been done; there is much more philosophical water that must flow under the bridge before we can be convinced that Papineau's suggestion has a solid intellectual basis.

Three Outstanding Issues

In this final section, I want to consider three important matters that flow from my defence of legal epistemology.

(a) How much Legal Epistemology do we need?

The argument has been that legal doxasticism is supported by the idea that courts should only attribute responsibility when members of the society they represent could reasonably do the same. This is why, contrary to ESF, I do not think that a preoccupation with rational belief in the criminal justice system amounts to a fetish. But there is a sense in which ESF may have a point. The point is this: it may *not* be necessary to engage in extremely technical epistemological jousting for the purposes

of criminal law epistemology. There may be a level of sophistication in our theorising—perhaps a level that we have already reached—that would suffice to make good on the argument of this paper.

Extant debates in legal epistemology, driven by issues arising from the use of statistical evidence in the courtroom, often turn on epistemic concepts that are rather complicated—e.g. sensitivity, normalcy, safety—to articulate with precision. It is not always clear the extent to which these concepts fully track the 'folk' notion of rational belief. Sometimes experimental philosophy is used to support certain views in epistemology, but philosophers do not typically regard such evidence as decisive or as a hard constraint on their theorising.

There is therefore a danger of legal-epistemological debate outstripping the folk notion of belief. Under the view discussed here, criminal judgements need to be able to command widespread confidence that only the morally responsible are being held to account.[29] If it turns out certain complex epistemic concepts diverge from our everyday concepts of rational belief, then, arguably, caring about these concepts would amount to a fetish in the way that ESF suggest. In a sense, according to my argument, it ought to be the folk notion of rational belief that provides the foundation for legal epistemology.[30]

(b) Shifty belief, shifty proof

Legal doxasticism claims that the standards for criminal proof must at least satisfy the requirements of rational belief. However, a controversial yet fairly popular view in contemporary epistemology is that the standards for rational belief vary with the stakes, which we can roughly conceptualise as the costs of error.[31] If legal doxasticism

---

[29] I should say, it certainly seems to be true that proponents of criminal law epistemology are correct that people *are* sceptical of convictions based on purely statistical evidence. For example, see Wells (1992), for influential empirical work.

[30] Of course, we might also be interested in idealised communities where every member possesses the best notion of rational belief. This is where we could see a divergence between ideal and non-ideal theoretic versions of legal epistemology.

[31] The literature here is voluminous, see Kim and McGrath (2019) for a recent collection. Moss (2021) articulates a version of this view regarding legal proof in terms of knowledge. Bolinger (2018) also defends a version of this view, although she is primarily concerned with the individual rather than legal case.

is well-founded, the truth of this stake-sensitivity thesis would have some rather important consequences for how we approach legal proof. Namely, it would mean that the minimum standards for criminal proof should vary with the stakes.[32]

At first blush, it may appear as if this possibility would reintroduce some of the worries we discussed earlier regarding the contingent costs of false acquittals—e.g. the variable dangers of recidivist crime and importance of incapacitating the dangerous—which would mean that doxasticist standard would vary unpredictably between different types of crime. However, this worry is misplaced, even if it were the case that members of society were attuned to the changing empirical facts about recidivism. The reason it is misplaced is that legal doxasticism is only a norm about guilty verdicts, not acquittals. So, the only costs relevant would be the costs of false convictions. In practice, this would mean something like the following: more serious crimes come with harsher punishments than less serious crimes, which means that the costs of error for the accused are higher, so we should use higher standards before convicting someone of the most serious crimes. [33] There is a sense in which this accords with intuition: it does seem plausible that we should demand more certainty to convict of murder than petty theft, given the massively increased severity of punishment for the former. However, I have not here taken any stance on whether shifty views of the rationality of belief are correct, nor on whether these views reflect the folk conception of belief.[34] Nonetheless, their truth would have some very interesting upshots for the project of legal epistemology.

(c) Doxasticism: necessary, not sufficient

To close, it is worth bearing in mind that legal doxasticism is a necessary rather than a sufficient condition on criminal conviction. One can accept all I have said so far, but still maintain that the appropriate standards for criminal conviction ought to be *more*

---

[32] Although, see Picinali (2013) for an argument against the idea that we can plausibly interpret the 'reasonable doubt' standard in such a flexible way.

[33] I expand on this issue in Ross (ms).

[34] All empirical evidence about folk belief is controversial, but a recent large study by Rose et. al (2019) suggests knowledge ascriptions do not vary with the stakes. This, of course, may not generalise to other conditions on rational belief, even if the authors are right about knowledge.

demanding than those for rational belief. Indeed, this is a rather attractive idea.[35] There are various different routes to this conclusion. One type of argument would be principle-based: for example, a rights-based argument against mistaken conviction, or a contractualist argument about the types of procedures that we can rationally agree to.[36] If these arguments do indeed secure a standard of proof in every way more demanding than the standards of rational belief, then the idea that we must rationally believe in order to appropriately convict would be *theoretically overdetermined*. However, there may be other contingent arguments for having a standard of criminal proof more demanding than the standard of rational belief. For example, perhaps, contrary to what Laudan suggests, a cost-benefit analysis actually currently recommends a very high standard of criminal proof (e.g. if Laudan has grossly overestimated the level of criminal reoffending). Such a cost-benefit argument, as we have argued, would only be a contingent one, and is hostage to changing empirical circumstances. Doxasticism is still a powerful principle, even if cost-benefit analysis currently recommends a very demanding standard of proof—it provides a *'fail safe'*, on which no matter how the empirical situation may change, we must never fall below the standards of rational belief when deciding whether to convict someone of a crime.

Conclusion

This paper has clarified and defended criminal law epistemology, aiming to place it on a solid normative footing in response to a spate of recent challenges.

The core commitment of criminal law epistemology is 'legal doxasticism'. This is the idea that guilty verdicts are appropriate only if a full belief in guilt would be rational, given the admissible evidence. I have argued that the very nature of punishment—following an institutional inquiry into moral accountability—requires asking whether we believe that the subject is culpable for their wrongdoing. Otherwise, we simply treat those being judged like a vector of risk, and criminal justice ceases to be a distinctive institution. Even if corporate bodies like courts can diverge from individuals in how they utilise certain concepts, it is nevertheless a basic condition of legitimacy for

---

[35] Indeed, in practice the standards for conviction are higher than those of belief in light of the various evidential rules and safeguards that we find in the trial context. Various types of evidence are excluded from legal adjudication that an individual would be free to consider (e.g., hearsay evidence).

[36] For example, Li (2015) provides a contractualist account of criminal law.

criminal courts that their judgements (aim to) conform to the judgements that a reasonable member of society would make if they considered the case for themselves. This does not mean that all aspects of criminal law epistemology are beyond reproach. For example, some of the more technical examples of the genre may diverge too far from folk notions of rational belief. But the basic project, so I have claimed, is well-motivated.

I want to close with a very important caveat. Namely, it remains to be seen whether *civil law* epistemology has similarly robust foundations. Civil adjudication is constituted by different norms and serves an altogether different social function than the criminal law—there is no guarantee that just because (so you will think if you agree with me) criminal law epistemology is well-founded, that the same can be said of its civil law counterpart.[37] Given that a great deal of ink has been spilled over civil law epistemology over the last decade, this challenge is one that we would do well to tackle head-on.[38]

References.

Backes, Marvin. (2020). Epistemology and the law: why there is no epistemic mileage in legal cases. *Philosophical Studies*, 177, 2759–2778.

Bolinger (2018). The rational impermissibility of accepting (some) racial generalizations. *Synthese*, 197, 2415–2431.

Blome-Tillmann, Michael (2017). 'More Likely Than Not' - Knowledge First and the Role of Statistical Evidence in Courts of Law. In Adam Carter, Emma Gordon & Benjamin Jarvis (eds.), *Knowledge First - Approaches in Epistemology and Mind*. Oxford, UK: Oxford University Press. pp. 278–292

Buchak, Lara (2014). Belief, credence, and norms. *Philosophical Studies,* 169 (2): 1–27.

Di Bello, Marcello (2019). Trial by Statistics: Is a High Probability of Guilt Enough to Convict? *Mind* 128 (512):1045-1084.

Enoch, D., Spectre, L., & Fisher, T. (2012). Statistical evidence, sensitivity, and the legal value of knowledge. *Philosophy & Public Affairs*, 40(3), 197–224.

Enoch, D., Fisher, T., Spectre, L. (2021). Does legal epistemology rest on a mistake? On fetishism, two-tier system design, and conscientious fact-finding. *Philosophical Issues*, 1–19.

Epps, Daniel, (2015). The Consequences of Error in Criminal Justice. *Harvard Law Review*, Vol. 128, 4: 1065-1151

Fratantonio Giada (2021). Evidence, Risk, and Proof Paradoxes: Pessimism about the Epistemic Project. *The International Journal of Evidence & Proof* 25(4): 307-325.

---

[37] See Ross (2021b) for some arguments on the distinctiveness of civil law epistemology.
[38]

Gardiner, Georgi. (2017). In defence of reasonable doubt. *Journal of Applied Philosophy*, 34(2), 221-241.

Gardiner, Georgi (2019). The reasonable and the relevant: Legal standards of proof, *Philosophy & Public Affairs* 47/3: 288–318.

Günther, Mario (2021). Epistemic sensitivity and evidence, *Inquiry*, DOI: 10.1080/0020174X.2021.1936158

Kaplan, John. (1968). Decision Theory and the Factfinding Process. *Stanford Law Review,* 20: 1065-1092.

Kim, Brian & McGrath, Matthew (eds.) (2019). *Pragmatic Encroachment in Epistemology*. Routledge.

Laudan, Larry (2006). *Truth, Error, and Criminal Law: An Essay in Legal Epistemology*. Cambridge University Press.

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 *Oxford Studies in the Philosophy of Law – I.*

Li, Hon-Lam (2015). Contractualism and Punishment, *Criminal Justice Ethics*, 34:2, 177-209

Lillquist, Erik. (2002). Recasting reasonable doubt: Decision theory and the virtues of variability. *UC Davis Law Review,* 36: 85-197.

List, Christian (2006). The discursive dilemma and public reason. *Ethics* 116 (2):362-402.

List, C., & Pettit, P. (2002). Aggregating sets of judgments: An impossibility result. *Economics and Philosophy*, 18 , 89-1 10

Littlejohn, Clayton 2020. Truth, knowledge, and the standard of proof in criminal law, *Synthese* 197: 5253–5286.

Moss, S. (2018). *Probabilistic knowledge*. Oxford: Oxford University Press.

Moss, S. (2021b). Knowledge and legal proof. *Oxford Studies in Epistemology*, 7.

Papineau, D. The disvalue of knowledge. *Synthese* 198, 5311–5332 (2021).

Pardo, M. S. (2018). Safety vs. sensitivity: Possible worlds and the law of evidence. *Legal Theory*, 24(1), 50−75.

Pardo, M. S. (2019). The paradoxes of legal proof: A critical guide. *BUL Rev.*, 99, 233.

Pettit, Philip (2007). Responsibility Incorporated. *Ethics* 117, 2: 171–201.

Picinali, Federico (2013). "Two meanings of 'reasonableness': Dispelling the 'floating' reasonable doubt." *The Modern Law Review* 76.5: 845-875.

Pritchard, D. (2018). Legal risk, legal evidence and the arithmetic of criminal justice. *Jurisprudence*, 9(1), 108–119.

Redmayne, Mike (2008). Exploring the proof paradoxes. *Legal Theory* 14 (4):281-309.

Ribeiro, Gustavo. (2019). The Case for Varying Standards of Proof, *San Diego Law Review,* 56: 161-219.

Rose, D., Machery, E., Stich, S., Alai, M., Angelucci, A., Berniūnas, R., Buchtel, E.E., Chatterjee, A., Cheon, H., Cho, I.-R., Cohnitz, D., Cova, F., Dranseika, V., Lagos, Á.E., Ghadakpour, L., Grinberg, M., Hannikainen, I., Hashimoto, T., Horowitz, A., Hristova, E., Jraissati, Y., Kadreva, V., Karasawa, K., Kim, H., Kim, Y., Lee, M., Mauro, C., Mizumoto, M., Moruzzi, S., Olivola, C.Y., Ornelas, J., Osimani, B., Romero, C., Rosas Lopez, A., Sangoi, M., Sereni, A., Songhorian, S., Sousa, P., Struchiner, N., Tripodi,

V., Usui, N., Vázquez del Mercado, A., Volpe, G., Vosgerichian, H.A., Zhang, X. and Zhu, J. (2019), Nothing at Stake in Knowledge*. *Noûs*, 53: 224-247.

Ross, Lewis (2020). Recent work on the proof paradox. *Philosophy Compass*, 15(6), e12667.

Ross, Lewis (2021a). Rehabilitating Statistical Evidence. *Philosophy and Phenomenological Research*, 102(1): 3-23.

Ross, Lewis (2021b). Justice in epistemic gaps: The 'proof paradox' revisited. *Philosophical Issues*, 31: 315– 333. https://doi.org/10.1111/phis.12193

Ross, Lewis (2021c). Legal proof and statistical conjunctions. *Philosophical Studies* 178: 2021–2041 https://doi.org/10.1007/s11098-020-01521-z

Ross (ms). Criminal Proof: Fixed or Flexible?

Roth, Andrea (2010). Safety in Numbers - Deciding When DNA Alone is Enough to Convict, 85 *NYU Law Review* 1130.

Smith, Martin (2018). When does evidence suffice for conviction? *Mind* 127/508: 1193–1218.

Wells, G.L. (1992). Naked statistical evidence of liability: Is subjective probability enough? *Journal of Personality and Social Psychology*, 62, 739–752.