The disvalue of knowledge

David Papineau^{1,2}

Received: 11 February 2019 / Accepted: 21 September 2019 / Published online: 4 October 2019
© The Author(s) 2019

Abstract

I argue that the concept of knowledge is a relic of a bygone age, erroneously supposed to do no harm. I illustrate this claim by showing how a concern with knowledge distorts the use of statistical evidence in criminal courts, and then generalize the point to show that this concern hampers our enterprises across the board and not only in legal contexts.

Keywords Concept of knowledge · Statistical legal evidence · Norm of belief · Norm of action · Norm of assertion · Reactive attitudes

1 The prehistory of knowledge

Bertrand Russell held that we would be better off without the concept of causation. He said that it, “like much that passes muster among philosophers, is a relic of a bygone age, surviving, like the monarchy, only because it is erroneously supposed to do no harm” (1912, p. 1).

Russell was wrong about causation. But his strictures apply perfectly to something else that passes muster among philosophers—namely, the concept of knowledge. This is indeed a relic of a bygone age, and moreover one that does appreciable harm. Or so at least I shall argue.

Let me begin with a genealogical account of the origins of the concept of knowledge. Other philosophers have offered genealogical accounts of the concept (for example, Williams 1978; Craig 1990). But such accounts are generally offered as vindictory fables, merely possible scenarios intended to cast the concept of knowledge in a good light. By contrast, I am interested in actual origins, and will suggest that they should make us suspicious of the concept’s credentials.

✉ David Papineau
david.papineau@kcl.ac.uk

¹ King’s College London, London, UK

² Graduate Center, City University of New York, New York, USA

There is good reason to suppose that the concept of knowledge is both phylogenetically and ontogenetically primitive.

Many animals are able to discriminate between agents who are and are not acquainted with some fact, even when they are insensitive to the possibility of agents *misrepresenting* their circumstances (Tomasello and Moll 2015). The basic category they employ is one of *openness* to a fact. At its simplest, this is applied on the basis of lines of sight: agents are assumed to know about what lies in their direction of gaze when nothing is in the way. I take this discriminatory ability to provide a foundation for the concept of knowledge. Those who have this ability can divide agents into those who *know* some fact and those who are *ignorant* of it.

Human children advance beyond this primitive ability as they develop, and become able to understand that, among those who are not acquainted with some fact, there is a further difference between the *false believers* and the merely ignorant: false believers are not only unacquainted with their circumstances, they positively misrepresent them. This is a significant intellectual advance, given that the false believers can be expected to behave like the knowers, and so can usefully be classified with them from a predictive and explanatory point of view (though of course their actions won't necessarily meet with success). A hungry agent who falsely believes some cake is in the pantry can be expected to go there, just like the person who knows it (even though the false believer won't then get the cake).

Human children can articulate the difference between knowledge and ignorance soon after they are able to talk. They only come to display an explicit grasp of the possibility of false belief, however, between the ages of three and four (Wellman et al. 2001).¹

In line with this, I am inclined to regard the concept of knowledge as an ancient precursor of our modern repertoire of psychological concepts. I have no doubt that our stone-age *homo erectus* ancestors two million years ago had a working concept of knowledge. Like their other primate relatives, they could distinguish agents who were open to some fact from those who were ignorant of it. But I doubt that they yet had an explicit concept of belief that would have allowed them to make more sophisticated categorizations and predictions.

2 Knowledge and true belief

Once we do have the concept of belief to hand, it is not clear why we should continue to be interested in knowledge. Of course, we will still want to distinguish, among believers, those whose thinking is in line with the facts, and those whose isn't, for the actions of the former but not the latter will generally meet with success. But the notion of *true belief* would seem to serve this function just as well as *knowledge*. A true belief that some cake is in the pantry will ensure the satisfaction of hunger just as well as knowledge to the same effect.

¹ Recent research indicates that human infants as young as 15 months, and even some great apes, can display an implicit ability to anticipate the actions of false believers (Onishi and Baillargeon 2005; Krupenye et al. 2016). This does not affect the point that an articulated understanding of these categories only develops in human children much later.

(Moreover, once we have the concept of belief to hand, we can also think about *degrees of belief*, and about their conformity to relevant objective probabilities. This offers further levels of conceptual sophistication with which to manage our doxastic affairs. But let us leave such refinements until later. For now I will stick to belief *simpliciter*. I will return to degrees of belief in Sect. 10.)

If knowledge calls for more than true belief, as on an intuitive level it clearly does, then why should it matter to anything whether or not the extra requirements are satisfied? What extra pay-off is delivered by knowledge but not true belief? As we shall shortly see, plenty of philosophers have sought to identify some such knowledge-dividend, some way in which the possession of knowledge makes a positive difference. Still, such attempts to find ‘the value of knowledge’ have met with very limited success. It has proved remarkably difficult, to say the least, to show how knowledge delivers better consequences than true belief or might otherwise be of benefit to its possessors.²

My own view is that there is no such pay-off and so no need to be interested in knowledge. The concern with knowledge is a stone-age hangover. Knowledge differs from true belief in ways that derive from the primitive idea of perceptual openness to a fact. But there is no advantage in focusing on this archaic category. We would do just as well if we forgot about knowledge and concerned ourselves only with true belief. Now that the more sophisticated notions of belief and truth are available, we get no benefit from pursuing knowledge.

In fact, I shall be arguing that things are worse than this. Not only does a concern with knowledge do no good, it does an appreciable amount of harm. Somewhat paradoxically, I shall defend this claim by focusing on what I regard as the most plausible account of a distinctive pay-off delivered by knowledge. Clayton Littlejohn has recently argued that we cannot make sense of certain legal practices unless we recognise that they are aimed at ensuring we do not convict and punish when guilt cannot be *known* (Littlejohn 2017). I think he is entirely right about this. As we normally think, we regard punishment when guilt cannot be not known as unacceptable, even when guilt is overwhelmingly probable.

My reaction is that, even if this how we ordinarily think, we ought not to. We might find it natural to shape our practices in this way, but the world would be a better place if we did not. We are being seduced by archaic ways of thinking into procedures that positively hinder our attempts to punish the guilty and save the innocent. We need to stop thinking in terms of knowledge.

I shall focus on the legal example because it brings out the issues in a particularly clear way. But I take the point to be of far more general significance. The premium placed on knowledge does not only distort the workings of courts of law. It hampers our thinking across the board as well as in legal contexts. I shall discuss these more general implications in the last two sections of this paper.

² What about the value of true belief? In previous papers I have argued that this is derivative rather than fundamental (Papineau 1999, 2013). For present purposes, however, we can put this issue to one side, and take it as given that true belief is valuable.

3 A good cup of coffee

Philosophers have long sought to explain why knowledge is superior to true belief. In the *Meno*, Plato suggested that knowledge is ‘tied down’ in a way that belief is not. The content of this suggestion is not entirely clear, but one reading of it has recently been defended by Timothy Williamson, who suggests that knowledge is more likely to persist in the face of counter-evidence than mere true belief, and to that extent often offers a better explanation of why some agents persist in their enterprises (Williamson 2000, ch. 2.4). This suggestion has not won widespread support, however, with a number of critics retorting that there is no obvious reason why knowledge should generally be more resilient in this way than strongly held belief (cf: Kaplan 2003; Magnus and Cohen 2003).

A rather different approach to the value of knowledge, going back at least to Bernard Williams, observes that anybody who values true belief will of necessity value knowledge, since the only effective means to the end of true belief is to acquire one’s belief in some reliable way. If we understand knowledge as true belief with a reliable provenance, this then shows how a desire for true belief will automatically generate a desire for knowledge (Williams 1978, ch. 2).

An initial response to this line is that doesn’t succeed in showing that knowledge has any *pay-off* beyond that delivered by true belief. It simply makes knowledge a sort of *means* to true belief. By way of analogy, it’s true that, if you are after a good cup of coffee, you will want to use a reliable coffee machine. Still, a qualitatively identical cup of coffee luckily issuing from some unreliable brewing process will in itself be no less valuable. It will taste just as sweet however it was made. Similarly, the wisdom of forming your beliefs reliably doesn’t show that the products of this process have any value beyond the value delivered by any true beliefs (Zagzebski 2003).

Some philosophers persist and say that *achievement* is itself valuable. A good cup of coffee that results from your coffee-making proficiency is something to be proud of, in a way that one you produce happenstantially is not. The same arguably applies to a true belief that is a consequence of intellectual skill, rather than irresponsible guesswork. The intellectual achievement is admirable, so the argument goes, and to that extent knowledge is valuable (Greco 2010).

One challenge facing this line of thought is to make good the underlying assumption that the intuitive category of *knowledge* coincides with *admirably achieved true belief*. I am happy to agree that it is indeed an admirable achievement to arrive at true beliefs by exercising one’s intellectual skills. It is a further task show that all such achievements count as *knowledge*, and vice versa. There are plausible counter-examples in both directions³ (due, if you ask me, to the excessive emphasis that the intuitive concept of knowledge places on a clear *causal* link between facts known and knowers).

More fundamentally, it is doubtful that the appeal to achievement really avoids the original cup-of-coffee objection. Even if we admire those who form true beliefs by exercising their intellectual skills, the worth of these skills would still seem to derive from the value of true belief. After all, these skills are praiseworthy insofar as they are an effective route to this valuable end. So we still don’t seem have an account of how

³ See Pritchard et al. (2018, Sect. 3).

knowledge is better than true belief, not just as a way of reaching it, but in terms of some additional pay-off.

So far, I have queried whether the superior value of knowledge can be explained in terms of its being ‘tied down’ or its status as valuable achievement. Even if these two suggestions are inadequate, however, this does not necessarily rule out some other account of the extra value delivered by knowledge. Perhaps a more penetrating analysis of knowledge will explain its extra virtue. Or maybe knowledge will resist any analysis in terms of other categories, but even so we will be able to identify some feature that accounts for its superiority.

But I am sceptical. In my view the failure of the philosophical literature to identify any additional pay-off from knowledge is itself evidence that there is none. That is by far the most obvious explanation for the fact that inordinate amounts of philosophical effort have come up empty-handed.

Some philosophers hold that knowledge is important because it provides norms for the management of belief, or action, or assertion, or evidence. However, these claims are typically defended by brute appeals to intuition and linguistic data, without any independent explanation of the benefit brought by these supposed norms, which is why I have not discussed them in this section. I shall consider these claims in Sects. 10, 11 and 12 below.

It is worth emphasising that the position I am defending here is consistent with the view that knowledge is a primitive notion that cannot be analysed in other terms. I am very sympathetic to this possibility. As I explained, I view knowledge as ancient category with deep roots in human history. In line with this, I am inclined to suppose that children acquire the concept by being introduced to stereotypical examples and then proceeding to extend the concept in ways that they find cognitively natural. There is no reason to suppose that such a basic recognitional concept can be reconstructed out of more sophisticated notions like belief, truth, reliability and so on.

Still, that knowledge is unanalysable doesn’t mean it is valuable. Humans might use a direct recognitional ability to decide which people know which facts. But, even so, the category so identified might possess no special value, except insofar as it coincides with true belief.⁴

Perhaps it is also worth making clear that I am not against *knowledge* itself, as opposed to the idea that knowledge offers extra value and so is worth aiming for as such. I am happy to allow that the concept of knowledge is perfectly cogent and precise (whether primitive or analysable), and so that many people are determinately in states of knowledge. And I allow that such states will generally be a good thing, in that their possessors will certainly have true beliefs, and moreover (at least in typical cases) will have exercised admirable intellectual skills in acquiring them. So I am not at all against states of knowledge themselves.

My objection is only to the idea that knowledge possesses some distinctive extra virtue and should therefore be pursued as such. As I have said, everything in our

⁴ It has been put to me that in practice humans need to start with the notion of knowledge in order to develop the notion of belief, and connectedly that they are more facile with the notion of knowledge than with compounds built out of belief, truth and reliable provenance. I am happy to concede this, just as I would concede that in practice humans need to start with erroneous ‘folk physics’ before they can develop any more sophisticated understanding of physical processes.

lives would go just as well, indeed would go better, if we simply stopped focusing on knowledge and concerned ourselves only with true belief. Of course, if we did, some people would still end up in states that satisfied the requirements for knowledge and would therefore enjoy the benefits of true belief. I don't want to stop that happening. What I'm against is the harms caused by our explicitly aiming for knowledge.

4 Statistical legal evidence

Clayton Littlejohn's recent 2017 paper 'Truth, Knowledge, and the Standard of Proof in Criminal Law' aims to explain why criminal courts are disinclined to convict on the basis of purely statistical evidence.

Consider this example:

'100 prisoners are exercising in the prison yard. Suddenly 99 of them attack the guard, putting into action a plan that the 100th prisoner knew nothing about. The 100th prisoner played no role in the assault and could have done nothing to stop it. There is no further information that we can use to settle the question of any particular prisoner's involvement' (Littlejohn, p. 1, following Cohen 1977; Nesson 1979).

Imagine now that one of these 100 prisoners is in the dock, and that the court has the above information, and that there is no possibility of any further evidence relevant to the defendant's guilt.

Littlejohn takes it, and I agree, that there is a very strong intuition that it would be improper to convict the defendant. Our only evidence is that the defendant belongs to a group of 100 people, 99 of whom committed the crime. It does seem quite wrong to convict solely on such statistical evidence.⁵

But this intuition is very puzzling. Compare the above case with one where the evidence is that an eye-witness saw a defendant steal some goods from a shop. In that kind of case, we have no corresponding intuition that conviction would necessarily be wrong.

But eyewitnesses are notoriously unreliable. Of the 362 USA prisoners exonerated of serious crimes on the basis of later DNA evidence since 1989, 70% were convicted on eyewitness evidence.⁶ True, this statistic in itself does not show that eyewitness evidence generally leads to false convictions. To show that, we would need to know about the proportion of innocents in all eyewitness cases, and there are obvious barriers to gathering such statistics. Still, such evidence as there is suggests that the ratio of false to true eyewitness convictions is worrying high (Wells and Hasel 2007).

Of course, any reasonable judiciary system will respond to such evidence by tightening up on the standards required of eyewitnesses. For example, the USA Department of Justice has recently issued strong guidelines about photo arrays and the prompting of eyewitnesses (Yates 2017).

⁵ Not only are our intuitions against statistically-based convictions, but so, at first pass, is legal doctrine. The full legal situation, however, is complicated (see Enoch and Fisher 2015; Smith 2018, footnote 2), and so in this paper I shall restrict my attention to intuition rather than legal practice.

⁶ <https://www.innocenceproject.org/causes/eyewitness-misidentification/>.

Still, such heightened evidential standards are not going to remove the puzzling disanalogy between eyewitness and purely statistical evidence. This is because the anomaly isn't a just quantitative matter of lower reliability being tolerated with eyewitnesses than with statistical evidence. Rather the disanalogy is qualitative. No amount of unreliability is intuitively acceptable in the context of statistical evidence. Yet we have no inclination to apply such absolute standards to eyewitness evidence.

5 Bad rationales

Nearly all of those who have discussed the matter uphold the intuitive thought that it is improper to convict on purely statistical evidence. It has proved surprisingly difficult, however, to find any good rationale for this intuition.

Sometimes it is suggested that purely statistical evidence is in tension with the requirement that truth be proved 'beyond a reasonable doubt' (Tribe 1971). But this suggestion is puzzling. Legal experts tend to be cagey about fixing a number for 'reasonable doubt'. Occasionally figures like 95% are mentioned (Simon 1969). Still, whatever the appropriate degree of certainty, we have already seen that a failure to reach some quantitative threshold is not the problem with statistical evidence. As I said, the resistance to statistical evidence is qualitative not quantitative. Statistical evidence can make doubt far less reasonable than the level tolerated with eyewitness convictions, but this still doesn't seem to make it adequate. Even after the statistical evidence makes guilt 95% likely (or 99%, or 99.9%,...) we continue to feel that it is an improper basis for conviction.

Colyvan, Regan and Person (2001) suggest that the problem with statistical evidence is that it would 'make it a crime to belong to a reference class'. Their thought is that the defendant found guilty on statistical evidence would be convicted *for belonging to a class most of whom were guilty*, rather than for being guilty per se.

It is by no means obvious, however, that we need to view statistical convictions in this light. The obvious alternative is to take the membership of the relevant class, not as constituting guilt, but rather as *evidence* for such guilt. The defendant is being convicted and punished (we trust) *for attacking the guard*. The defendant's membership of a class of whom 99% attacked the guard is merely fallible but strong evidence of this guilt. Moreover, there can't be anything in itself wrong with courts using such less-than-infallible tests for guilt, since they happily use such tests elsewhere, as for example with eyewitnesses.

So Colyvan et al. cannot take the inconclusivity of the statistical evidence to imply that we are condemning the accused for displaying the evidence, rather than committing the crime that this is evidence of. If they were right, we could equally object that someone convicted on eyewitness evidence was being convicted for being *said to have been seen robbing the shop*, rather than (we trust) for robbing the shop.

Enoch et al. (2012) and Enoch and Fisher (2015) argue that statistical evidence is not sufficiently *sensitive* to whether or not the defendant is guilty, and as a result does

not appropriately *incentivize* potential offenders to avoid transgressing. Their thought is that, if the defendant had known that conviction was assured on statistical grounds, whether or not they attacked the guard, then the defendant would have had no incentive not to attack the guard. By ensuring that conviction is appropriately tied to whether or not the particular defendant committed the crime, the legal prohibition on statistical evidence thus gets the incentives right.

This is an interesting suggestion. The idea that convictions should be appropriately sensitive to the defendant's guilt seems close to the heart of the intuition against statistical evidence, and I shall return to it briefly below. Moreover, it does seem reasonable to design legal practices, not just to reach factually accurate decisions, but also to provide incentives not to commit crimes. Still, the supposed advantage in terms of disincentives does not offer a convincing rationale for the general legal rejection of statistical evidence.

This is because the outlawing of statistical evidence provides extra disincentives only in very specific contexts. To see this, consider the prison-yard scenario once more. A no-statistical-evidence policy will only create an extra incentive not to attack if there *won't* be eyewitnesses (or other individually specific evidence) if you *don't* attack, but there *will* be such evidence if you *do* attack. After all, if you were an innocent prisoner, then you wouldn't be at risk, even given a statistical-evidence policy, if there were eyewitness evidence that you *hadn't* attacked—for that individual evidence would trump the statistical case against you. Conversely, even a no-statistical-evidence policy wouldn't encourage you not to attack unless there *would be* individualized evidence if you *did* attack—otherwise conviction *still* wouldn't be sensitive to whether you attacked or not, even under a no-statistical-evidence policy. In short, a very specific set of circumstances is required for a no-statistical-evidence policy to incentivize lawful behaviour. Enoch et al. (2012) contend that their account in terms of extra incentives fits legal practice across a range of contexts, but they do not explain why we should expect such extra incentives to be at all common.

6 Rejecting the intuition

Let me put me cards on the table. I don't think the rejection of statistical evidence is a good thing. I recognize there is a widespread intuitive feeling that purely statistical convictions are improper. But I say that we should ignore this feeling and allow the courts to use statistical evidence.

After all, if we want to convict the guilty and free the innocent, then it *can't* be a good idea to convict when the evidence before the court makes guilt $n\%$ likely (because we have an eyewitness) yet acquit when the evidence makes guilt even more probable (because the evidence is only statistical). To see this, note that there must be some probability-of-guilt threshold for convictions that maximizes the expected value of convicting-guilty and freeing-innocent minus the negative value of convicting-innocent and freeing-guilty. If you're sometimes convicting when guilt is $n\%$ probable,

yet freeing when guilt is even more likely, then you can’t be adhering to that threshold, and so can’t be maximizing the expectation of valued outcomes.⁷

Despite the simplicity of this argument, few have been prepared to reject the ban on purely statistical evidence. (Though see Lempert 1977; Laudan 2006, Steele MS.) Even in the absence of any agreed justification for the ban, most commentators remain convinced that some as yet unknown rationale must lie behind the strong intuition against statistical evidence. I now propose to undermine this conviction by identifying the real source of the intuitive resistance to statistical evidence and showing that it does not constitute any justification for the ban.

7 Guilt and lotteries

According to Littlejohn, the resistance to statistical convictions stems from the simple thought that it is improper to convict someone if we don’t *know* that they are guilty (even if they are and the evidence shows that’s very likely).

Littlejohn’s central idea is that knowledge of guilt cannot be derived from purely statistical evidence. Just as knowing that a lottery ticket is almost certain to lose doesn’t mean that we know it will lose, so knowing that the defendant is almost certain to be guilty on statistical grounds doesn’t mean that we know the defendant is guilty. And since we can’t convict if we don’t know, argues Littlejohn, statistical evidence can never justify conviction.

By contrast, according to Littlejohn, in the eyewitness case we are quite likely to know that they are guilty. Sure, eyewitnesses might only be 95% reliable, say, and so it is not ruled out that we don’t know. Still, in the 95% of good cases, where the eyewitness was indeed in visual contact with the crime, and transmitted this contact to us via testimony, then we will know, and can legitimately convict. The high reliability of eyewitnesses can put it beyond reasonable doubt that we know, and so can justify conviction.

In short, purely statistical evidence implies that we definitely don’t know, and so rules out conviction. Eyewitness and similar individually specific evidence can make it likely beyond reasonable doubt that we do know, and so can make it proper to convict.

I have no doubt that this is the right account of our intuitions about statistical evidence. It is an interesting question exactly why judgements deriving from eyewitnesses

⁷ Imagine the values of the different outcomes are as below and p is the probability of guilt. Then you should convict when $p - 18(1 - p) > -p$, that is when $p > 0.9$. More generally, it is clear that any assignment of values to the four outcome fixes a probability such that it will maximize expected value to convict just in case guilt is higher than that probability.

	Guilty	Innocent
Convict	1	− 18
Free	− 1	0

qualify as potential knowledge, whereas those deriving from statistical evidence do not, and I shall return to this issue in the section after next. But there is little doubt that the notion of knowledge does so discriminate. As I said at the beginning, the central motif in the idea of knowledge is *openness* to the fact known, courtesy of some unhindered causal path from the fact to the knower. In the case of eyewitnesses the putative causal path is clear, and so we don't hesitate to allow that we might have knowledge. But statistical evidence does not forge any causal connection between putative knowers and the target fact, except insofar as it is one among many other similar facts, and so strikes us a paradigm of the kind of evidence that cannot deliver knowledge.

8 Caring about knowledge

While I am very happy to agree with Littlejohn that a concern with knowledge *explains* our intuitive resistance to statistical convictions, I don't at all think that it *justifies* that resistance. As I see it, the concern with knowledge is preventing us from arranging our legal procedures in an optimal way. It is hindering our attempts to convict the guilty and free the innocent.

From my perspective, our intuitions are being driven by archaic ways of thinking. We have this primitive idea of being open to the fact known, and we are allowing it to condition our legal procedures. We want to be in causal contact with the facts, and feel we lack such contact when our evidence is statistical. And so we judge that statistical evidence can't possibly give us the kind of acquaintance with the facts that is needed for a legitimate conviction.

But that is all bad thinking, driven by the crude prehistoric notion of openness to the facts. After all, beliefs based on statistical evidence can be just as true as potential knowledge, and in many cases are far more likely to be so. Given this, we'd be better off, with respect to everything else that matters, if we freed ourselves from our prejudice against them. There is no real reason to hold that convictions based on knowledge are better.

It is worth noting that Littlejohn's account of the resistance to statistical evidence implies a kind of answer to our earlier puzzle about the *value* of knowledge, though of a surprising sort. We need to be concerned about knowledge, his view implies, precisely so that we can pursue the goal of *convicting when we know* and avoid the mistake of *convicting when we don't know*. The outcomes that we value in courts of law are partially constituted by states of knowledge.

In the last section, in making the case for statistical evidence, I took it that court procedures were aimed at generating true verdicts and avoiding false ones. From Littlejohn's perspective, this misunderstands the outcomes the courts aiming to achieve. They aren't concerned with truth and falsity per se, but rather with having our verdicts

accompanied by knowledge. It's good to convict when you know, and bad to convict when you don't.⁸

This seems an odd place to find a value for knowledge. Earlier we assumed that, if knowledge were indeed valuable, then its value would lie in some extra dividend, in some function served by knowledge but not by true belief. But the idea now before us is in effect that *knowledge is its own dividend*. Knowledge doesn't matter because it is good for something else, but because it is good in itself—or, more accurately, it matters because it is good in itself that knowledge be present and not absent when we convict.

However, precisely because it identifies no further function for knowledge, this kind of account of the value of knowledge is hard-placed to respond to those, like myself, who simply deny any such value. We sceptics can concede that ordinary people *desire* that convictions be accompanied by knowledge, and that this desire shapes the practice of the courts. But we think that this desire is misplaced. There is no value in its satisfaction. It is an upshot of a stone-age mode of thinking that does no good and only directs our efforts away from what it truly valuable.

9 Drawing the line

Let me return to the issue flagged at the end of the section before last. Suppose we agree that, at an intuitive level, convictions require knowledge. But how exactly does this work in practice in the legal context? It seems clear enough that eyewitness testimony is in, and prison-yard-type statistical evidence is out. But there are less clear-cut cases. What about DNA traces, records of previous convictions, plausible motives, and so on? Intuition is not always decisive on whether or not these sources of evidence are tarred with the same brush as purely statistical evidence.

Some of the literature on statistical evidence and the law takes up the challenge of drawing a principled line between eyewitness testimony and purely statistical evidence. For example, Judith Jarvis Thompson (1986) maintains that good evidence needs to be *individualized* and suggests that this requires some *causal connection* to the target fact. Enoch et al. (2012) and Enoch and Fisher (2015), as we have seen, maintain that a guilty verdict needs to be appropriately *sensitive* to the target fact. And Martin Smith (2018) argues that evidence needs to *normically* support the verdict, in the sense that it would be *abnormal* for the verdict to be mistaken given that evidence.

While these writers all note that their suggested requirements bear some similarity to the requirements for knowledge, they don't commit themselves to Littlejohn's thesis

⁸ So Littlejohn's decision table would look like this:

	Known Guilt	Unknown Guilt	Innocent
Convict	1	– 18	– 18
Free	– 1	0	0

and we will never convict on statistical evidence alone since in such cases we can be sure guilt is not known (cf: Littlejohn op cit Sect. 7).

that the courts are explicitly concerned with knowledge as such. Their aim is rather to understand the underlying criteria the courts use to decide the propriety of evidence and the rationale for these criteria.

From the point of view of my analysis so far, it is interesting to note how all these suggestions run into problems drawing the line in the right place. Thomson's causal requirement can only avoid counterexamples by becoming unacceptably vague (cf: Enoch et al. 2012, Sect. 6.2). The requirement of sensitivity fails to rule out all cases of statistical evidence (cf: Smith 2018, Sect. 2). And Smith's appeal to normalcy is hostage to the challenge of developing a precise account of what constitutes abnormality.

It is also interesting to note how none of these suggestions generates a plausible rationale for why we should convict only when evidence is on one side of the line. Thompson suggests that only individualized evidence can offer an appropriate 'guarantee'; the unreliability of eyewitness testimony, however, makes it unclear what this guarantee amounts to. Enoch et al. and Enoch and Fisher argue that sensitivity ensures that the law offers appropriate incentives; we have already cast doubt on this as a rationale for the overall legal ban on statistical evidence. As to Smith, he avoids committing himself to any account of why it is a good idea to require normalcy for conviction.

If Littlejohn's account of intuitive legal practice is right, these difficulties are not surprising. As Littlejohn tells it, our intuitions are being driven by the difference between cases where we can and can't *know* about guilt. This implies that explaining this difference in other terms would be tantamount to analyzing the concept of knowledge. In effect, Thompson and the others are trying to specify necessary and sufficient conditions for knowledge, albeit restricted to the legal context. Given the miserable track record of attempts at such explicit analyses of knowledge, it is scarcely surprising that this legal version of the analytic enterprise should run into trouble.

Again, if Littlejohn is right, no account of the line between proper and improper evidence is going to uncover any independent rationale for favouring the former over the latter. If combining conviction with knowledge is intuitively its own reward, as Littlejohn maintains, then there will be nothing more to be said in justification of our existing practice. It may be approximately true that we require convictions to be based on individualized, or sensitive, or normic evidence. But there is no further reason for this requirement, apart from that's how we like it.

10 Flat-out belief

Perhaps we need to think a bit more about the mental attitude required to condemn a defendant on the basis of statistical evidence alone. If you simply thought that there is a 99% chance that the defendant committed the crime, then would you really be in a position to go ahead with conviction? You might be able to say that your reason (you trust) is *that the defendant did it* (as opposed to this being 99% probable). But, even so, would your probabilistic attitude be an adequate basis for conviction? We might feel that conviction requires us to regard the guilt as *settled*, and not merely as highly probable.

After all, the judiciary system is now going to sentence the defendant to punishment, and to close down any further debate on the question of guilt (unless significant new evidence comes to light). To this extent, it will no longer regard the issue of guilt as open. And we might well feel that this calls for a *flat-out belief* that the defendant is guilty, as opposed to merely believing that this is highly probable.

Note that there is a significant gulf between flat-out believing some p and believing that p is highly probable. We can think of the latter probabilistic belief either as (a) a flat-out belief *that p is highly probable*, or (b) as a high degree of belief *that p* . But either way the probabilistic belief is very different from the flat-out belief that p . Most obviously, the flat-out belief will be *mistaken* if p is false (if the defendant is not guilty), but not so for the probabilistic belief. It can be quite right to hold *that p is highly probable*, or to correspondingly have a high degree of belief *that p* , even in cases where p is false. (After all, it was true that the defendant's guilt was 99% probable, even in the one-in-a-hundred case where the defendant is innocent).

If flat-out belief is indeed required for conviction, on the grounds that conviction requires the courts to regard guilt as settled, then this does provide a sort of rationale for the requirement that courts can't convict without knowledge. This is part of Littlejohn's view. Along with a number of other philosophers, Littlejohn favours a 'knowledge norm for flat-out belief': he thinks that flat-out belief is only warranted in cases of knowledge. Since statistical court evidence is never sufficient for knowledge, it precludes flat-out belief in guilt, and thus blocks the legitimacy of convicting on statistical grounds.

A first issue here is whether flat-out beliefs ever play a positive role in our intellectual economy. It is not obvious that flat-out beliefs are really needed for legal procedures, or indeed for any other intellectually guided activities. Probabilistic beliefs are far more sensitive instruments, allowing us to weigh possible outcomes according to their likelihood, and to choose our actions accordingly. At first pass, there seems every reason to suppose that dispensing with flat-out beliefs in favour of probabilistic beliefs across the board will only improve our decision-making.

On the other side is the thought that it is often pragmatically necessary to regard some issue as settled and not open to further debate. Humans with our limited cognitive powers can find it difficult to reach definite practical or theoretical conclusions if crucial premises in our reasoning only have a provisional status. In many cases, so the thought goes, we can only make progress if we place certain claims beyond the scope of further deliberation (cf: Harman 1986).

This is not the place to resolve this issue. I myself do not think that the pragmatic case for flat-out beliefs goes very deep. It seems to me that it only has a grip in cases where we are forced to decide things in a hurry. Still, I recognize that this is a complex issue that deserves fuller analysis. Fortunately, we can by-pass it here, as it is not decisive for the legal status of statistical evidence. Suppose that I concede, for the sake of the argument, that flat-out belief does sometimes play a necessary role in our cognitive economy, and in particular that it is called for if we are properly to convict criminals. It still doesn't follow that it is always improper to convict on statistical evidence.

This is because the 'knowledge norm' for flat-out belief is highly questionable. This is the principle that is supposed to block our forming flat-out beliefs on the basis

of statistical evidence. According to the ‘knowledge norm’, since statistical evidence doesn’t suffice for knowledge, it can’t warrant flat-out belief. But there is a perfectly sensible alternative to the knowledge norm for flat-out belief—namely, the ‘Lockean’ principle that you can form a flat-out belief once the evidence raises the probability of some claim beyond some appropriate threshold. If we adopt this principle, instead of a ‘knowledge norm’, then the courts will be able to form flat-out beliefs on the basis of statistical evidence whenever that evidence raises the probability of guilt sufficiently high, and so in such cases will be able to regard guilt as settled in a manner appropriate to conviction and sentence. This then opens the way once more to pursuing accurate verdicts without the distractions occasioned by any misplaced concern with knowledge.

True, any such Lockean principle will inevitably threaten inconsistency in our flat-out beliefs. If we flat-out believe that one defendant is guilty, on the statistical evidence that 99 of the 100 committed a crime, then what’s to stop us adopting the same attitude to all the other members of the group, and thus ending up flat-out believing of each that they are guilty, even though we already flat-out believe that only 99 of them are?

But defenders of a knowledge norm are ill-placed to object to the Lockean principle on these grounds. This is because their management of flat-out beliefs also leads to inconsistencies. As has been pointed out before, while a knowledge norm might block lottery paradoxes, it still leaves us open to preface paradoxes.

Littlejohn helpfully transposes the point into the legal context. Suppose 100 prisoners have been convicted in 100 trials, each on the basis of eyewitness testimony. After all the trials are over, ‘a perfectly reliable observer tells you that precisely one of these people had been framed. Alas, your informant dies before he can identify the person’ (Littlejohn 2017).

Even after you know about the framing, the knowledge norm will still lead you to flat-out believe that each of these prisoners is guilty. This is because, if the knowledge norm is to be of any practical use, it needs to be accompanied by the prescription that you should flat-out believe *if it is sufficiently likely* that you know. Requiring absolute certainty that you know would stop you ever forming any flat-out beliefs, and in particular would stop you ever convicting on eyewitness evidence. Since in each case it is 99% likely that you know the prisoner is guilty, courtesy of good eyewitness evidence, you should therefore have 100 flat-out beliefs. Yet at the same time you should clearly also flat-out believe that one of the defendants is innocent, since a perfectly reliable informant told you so.

So the knowledge norm also gives rise to inconsistencies in flat-out beliefs. Given this, it would seem unreasonable to condemn Lockeanism for inconsistency. No doubt the friends of knowledge will feel that the inconsistencies resulting from their norm are less malign than those stemming from Lockeanism. But it is hard to see what argument could back up this feeling, except one that already assumed that flat-out belief without knowledge is unacceptable. (If you ask me, inconsistencies are simply the price we must pay for wielding the crude intellectual instrument of flat-out belief, whatever norm we govern it by.)

11 Personal relations

Littlejohn doesn't only hold that we shouldn't convict people that we don't know to be guilty. He also maintains that we shouldn't *blame* people that we don't know to be guilty.

This moves us out of the legal context and into the realm of personal reactions. This places some of the issues in a new light.

Imagine that you notice certain changes in your spouse's behaviour—frequent visits to the hairdresser, working into the evening, and so on. You read in a reputable psychology magazine that 99% of people behaving in this way are having an affair.

Would it now be all right for you to *blame* your spouse for infidelity? Intuitively, I take it, that would be wrong, almost pathological. You aren't in any position to blame your spouse. You can't cogently blame someone if your evidence is solely that they belong to a category 99% of whom have transgressed.

But suppose instead that a friend reports seeing your spouse canoodling with someone in a restaurant. Now it doesn't seem nearly so crazy to blame your spouse for infidelity. And this intuition stands even if we all know that the chance of malicious misinformation in such cases is much higher than 1%. (Of course, if your spouse is innocent, then it will not be easy to make amends. But blaming on the basis of your friend's word does not seem outlandish, in the way that blaming on statistical evidence does.)

Our intuitions in these cases precisely mirror the intuitions in the legal context. Intuitively we can't properly blame our spouse on statistical evidence, just as intuitively we can't properly convict a defendant on statistical evidence. However, if the evidence is eyewitness testimony, our intuitions go the other way, and it seems reasonable both to blame our spouse and to convict the defendant—even if the likelihood of error is higher than it would be on statistical evidence.

I am persuaded, following Littlejohn, that these parallel intuitions have a parallel explanation. With both blame and conviction, we allow eyewitness but not statistical evidence because we intuitively take both responses to require knowledge. Since statistical evidence alone cannot deliver knowledge, we can neither blame nor convict on statistical grounds alone.

Given this parallel, one option would be to respond as I did in the legal case, and urge that we 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.

But there is more to be said here. In 'Freedom and Resentment' (1962) P. F. Strawson argued that the 'reactive attitudes' are too central to our lives for it to make sense to try to repress them on theoretical grounds. Strawson's reactive attitudes include blame, along with praise, resentment, gratitude, guilt, remorse and other such natural emotional responses. In Strawson's view, it would clearly be undesirable, and in any case practically impossible, for us to stop reacting to each other in these ways, and so no abstract philosophical argument could possibly show that we ought to.

Strawson does not talk about knowledge in relation to the reactive attitudes. But if knowledge is intuitively connected to blame in the way indicated above (and similarly

to praise, resentment, and so on), then these connections might plausibly be validated along with Strawson's more general defence of the reactive attitudes. Perhaps it is a central part of our social fabric that we owe loyalty to those close to us, and perhaps part of so being loyal is not to blame (or praise, or resent...) them for things we don't know they did. This argues that it might do undue damage to our lives if we were to try to disconnect blame and similar reactive attitudes from knowledge.

This line of thought raises a number of interesting questions. How essential to our lives is the particular connection between knowledge and the reactive attitudes? It would be interesting to know whether the intuitive inadmissibility of blaming on purely statistical evidence is a human universal, rather than something peculiar to modern western societies.

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 (perhaps because of a recognition that the concept of knowledge is a relic of a bygone age...). 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. Strawson himself allows that there is often good reason to modify or redirect the 'web of human attitudes' at a piecemeal rather than a wholesale level. (Note in particular that we could still preserve the requirement that blame calls for flat-out belief and not just high credence, if this seemed necessary or desirable, even after loosening the tie to knowledge. For we could always come to condition our flat-out beliefs on high probability rather than knowledge, in the Lockean manner outlined in the last section.)

I have no strong views at this point. To the extent we are simply concerned to blame the guilty and not the innocent, then we will certainly do better to ignore issues of knowledge, for all the reasons rehearsed so far in this paper. But I see that there might be a cost. Perhaps it would indeed impoverish our lives if we started to blame people in our social circle, and adopt similar reactive attitudes, on purely statistical evidence, even in cases when we don't know that they transgressed. Maybe this would dilute the kind of loyalty we can expect of those close to us, and thereby render our social world less congenial.

I rather doubt that this cost would be significant.⁹ Still, rather than pursue the point, suppose that I do now concede that something of value would be lost if we started blaming people on statistical grounds. This would require me to qualify my strong thesis that the concept of knowledge does no good and only harm. For I would now be agreeing that knowledge can make a positive contribution after all—not because it delivers some further dividend, but because it is worthwhile in itself that we adopt certain reactive attitudes only on the basis of knowledge and not on merely statistical evidence.

⁹ Note that in practice there will nearly always be other reasons for not blaming etc on statistical evidence alone, apart from any requirement that such reactions demand knowledge. For a start, we will normally have much richer information about individuals than that they are members of some statistical reference class, which will block any direct inference from reference class probabilities to judgements about individuals. Moreover, even when we don't have such extra information, we will often owe it to the individuals concerned to acquire some before forming any firm conclusions.

Still, note what a small and curious role this now assigns to the category of knowledge. Far from playing a central role in the logic of enquiry, knowledge turns out to matter only to the ethics of interpersonal relationships.

Most philosophers assume that the category of knowledge is crucial to the overall management of our intellectual lives, by constituting the aim of enquiry, or some such. But we have found nothing to back up that assumption. We would seem able to manage our intellectual affairs just as well, indeed better, by adopting whatever strategies are best designed to lead us to true beliefs, without bringing knowledge into it. The category of knowledge turns out to be significant only as a very specific device to stop us too readily blaming (praising, resenting...) those close to us.¹⁰

12 Knowledge norms

If knowledge plays a positive role in structuring our emotional lives, then perhaps it plays similar roles in other places too. I am doubtful. As I have said, I view the concept of knowledge as an archaic hangover. In light of this, I am prepared to accept that it might play an ineliminable role in structuring our evolved fabric of social relations. But I think it unlikely that it makes any positive contributions beyond this.

Some philosophers hold that knowledge provides the norm for rationally based *action*. As they see it, it is rational to act on some reason only if you know it (Hawthorne 2004; Stanley 2005; Hawthorne and Stanley 2008; Smithies 2012). If they were right, this would certainly show that knowledge is generally of value, since knowledge would be a prerequisite for any rational action.

The thesis that knowledge provides the norm for action is a generalization of the claims considered in this paper so far, that you need knowledge to convict criminals, or to react with blame or similar attitudes. Given this, it is unsurprisingly open to similar objections.

Suppose your ticket in a lottery has a one-in-a-million chance of winning the £1,000,000 prize. In the absence of any further information, it would intuitively be irrational to throw your ticket away simply on the grounds that it's not going to win. But if you read in the newspaper that some other ticket has won, then this action intuitively becomes rational.

The natural explanation for this asymmetry is that you shouldn't treat something as a reason for action if you don't know it. You didn't know your ticket wouldn't win when your evidence was merely statistical, but you almost certainly came to know this when you read the newspaper report.

But, as before, allowing knowledge to guide our actions in this way can lead us to act in ways that are manifestly not in our interests.

¹⁰ Strawson talks of the 'generalized' and 'vicarious' reactive attitudes that we adopt towards people we do not know. These too arguably play an essential role in a healthy society, and perhaps in its legal institutions in particular. Still, it is a further question whether these generalized reactions call for *knowledge* of wrongdoing and not just good statistical evidence. Once we move away from the special ties of loyalty involved in close personal relationships, we lose the Strawsonian rationale for preserving the knowledge requirement. Nothing would seem to be lost, and much gained, if our generalized reactive attitudes and any associated flat-out beliefs were responsive not just to causally-based knowledge but also to often more reliable statistical evidence.

Suppose that the newspaper in which you read the winning ticket number does occasionally make mistakes, and moreover that the number it prints shares the first three digits with yours. In itself, this doesn't of course mean that you haven't gained knowledge from the paper that your ticket lost. It's still very likely that you have, and if the knowledge norm for action has any practical implications, it will say you should now bin your ticket. Still, the probability of your ticket winning, after your reading the paper, could well now be far *higher* that it was when your only evidence was statistical. (Maybe 1% of the time the copy-takers lose concentration after the first three digits.) Given this, how can it possibly be a good idea to hang onto your ticket before you read the newspaper, but throw it away afterwards?

We can imagine a variant of this example where you are offered a choice between two tickets, a normal one with a one-in-a-million chance, and one that shares the first three digits with the winning number reported in the less-than-fully-reliable newspaper. Relative to your epistemic situation, let us suppose, the second is ten times more likely to win. But the knowledge norm of action implies that you should prefer the first ticket to the second. This seems little short of perverse. Maybe the practicalities of human decision-making mean that you will do well to forget about your ticket winning once you read the newspaper. But, if so, how could the same practicalities not advise you to do the same with a normal ticket that is far less likely to win?¹¹

Let me be clear about my argumentative stance. I am not disputing the intuitive plausibility of the knowledge norm for action. I am quite happy to allow that it is line with everyday intuition to act on the newspaper report but not on mere improbability (just as everyday intuition views it as legitimate to convict on the basis of eyewitness but not statistical evidence). My position is that these are bad intuitions. They are leading us to act in ways that are not in our interests, and so should be resisted.

Defenders of the knowledge norm for action typically rest their case entirely on everyday intuitions. They do not point to some advantage that the norm delivers, some way in which it makes our lives go better. They just appeal to what ordinary people would normally say. As such, their arguments simply fail to engage with my claim that the norm is a bad thing.¹²

There remains the Littlejohn-style option of arguing that it is valuable in itself that we should act on knowledge. It's not that doing this brings some further benefit. It is its own reward. It's good per se to act on knowledge, but not on merely statistical evidence.

I'm not sure how far this further thought retains the support of intuition once we move past the specialized demands of the law and the reactive attitudes. I would have thought that most people thinking about lottery tickets see the value in money (and in not taking unnecessary pains to hang on to their tickets) but are quite unconcerned

¹¹ Some hold that *whether x knows p* is relative to the context of assessment and/or to how much is at stake for x with respect to p. The point I am making is insensitive to these complications. The knowledge norm for action will still favour knowledge over often more reliable statistical data even after we fix the context of assessment and the stakes.

¹² Some hold that we need to *know* some fact in order to *act in the light of* it, or *be guided by* it, or even to *act for* it (Unger 1975; Hyman 1999, 2011). Perhaps that is right. My response is then that there is no virtue in basing decisions only on facts to which we bear such putatively knowledge-requiring relations. Again, such a requirement may have the support of intuition, but no good comes of it.

about the meta-issue of whether their choices are guided by knowledge. Still, even if they did feel the intuitive pull of this generalized meta-concern, I would say they are misguided. They would be seeing value where there is none. An archaic mode of thought would be seducing them into wanting something that is of no worth.

In Sect. 10 above I discussed a knowledge norm for *flat-out belief*. This is closely related to a knowledge norm for action and faces the same difficulties. In Sect. 10 I observed that a Lockean norm of flat-out belief would do better at convicting the guilty and freeing the innocent. The point generalizes. From the point of view of achieving any of our ends, a knowledge norm for flat-out belief must do worse than a Lockean one, given the way it favours causally-based knowledge over often more reliable statistical evidence.

As it happens, it is debatable whether a knowledge norm for belief commands the same intuitive support as one for action. Everyday talk of ‘belief’ does not focus unequivocally on flat-out belief, as opposed to various weaker doxastic attitudes, and this complicates the issues (Hawthorne et al. 2016). In any case, the intuitive case for a knowledge norm of flat-out belief is not my concern. Even if this norm had the full support of everyday intuition, I would still say that it is a bad idea.

Let me now turn to the idea that knowledge is the norm of *assertion*. According to this norm (defended by Williamson 1996, 2000; DeRose 2002) you should only assert something if you know it.

Here too the standard defence invokes intuitions about lottery-style examples. If your information is solely that your ticket is a million-to-one against, it seems wrong to say outright that ‘My ticket isn’t the winner’. But this becomes intuitively acceptable once you read in a newspaper that another number is the winner (cf: Williamson 1996, pp. 498–505).

This norm is open to the same kind of objections as the other knowledge norms. If assertion is a means of transmitting information for others to rely on, how can it be a good idea always to favour causally-based knowledge over statistically based claims that are often more likely to be true?

Once more, intuition offers no case against the view that we would be better off if we didn’t require assertions to be backed by knowledge. This practice might seem natural to everyday thinking, but that doesn’t establish that it’s benign.

On some views, the practice of assertion is *constituted* by the norm of knowledge, and so there’s no question of it not being governed by this norm: someone who was not sensitive to the norm simply wouldn’t be making assertions (Williamson 1996). But this only moves the issue around. The question is now why we should go in for assertion, rather than engage in some other communicative practice, and once more intuition alone is impotent to supply an answer.

What communicative practice might we engage in, if not one governed by a norm of knowledge? This is not the place to resolve this issue, but it is worth observing that many languages require speakers to mark all indicative utterances with an *evidential* that indicates the provenance of their claim (such as first-hand observation, testimony, inference, and so on, depending on the range of evidentials available) (Aikhenvald 2004). This seems an improvement on a system that simply assumes that all indicative utterances are geared to a single standard. This is not to argue that all languages need a formal system of evidentials. Speakers can always use other means to con-

vey what kind of backing their claims enjoy. Indeed it is hard to believe that this isn't already demanded of speakers of non-evidential languages, in any context where hearers are likely to rely on their claims. The idea that communication hinges on a uniform category of 'assertion' answering to some undifferentiated standard strikes me as something of a philosophical myth.

Finally, let me mention the idea that the notion of knowledge is needed to understand *evidence*. Timothy Williamson has argued that our evidence always coincides with we know and defends this equation as the best explanation of various intuitive facets of evidence.

This topic deserves fuller discussion, but let me make some quick remarks. A first question is whether we need a uniform notion of evidence, rather than a more flexible understanding of legitimate starting points for different kinds of investigative reasoning (Joyce 2004). Second, even if we do need such a uniform notion, why should it require flat-out belief, as opposed to credences on which we can Jeffrey-conditionalize (Jeffrey 1965, ch. 11)? Third, even if these two points are granted, it is debatable whether the intuitive structure of evidence is best explained by equating it with knowledge rather than with justified non-inferential belief, or true non-inferential belief, or perhaps true justified non-inferential belief; Alvin Goldman has argued in detail that the points Williamson appeals to are equally well accommodated by the first of the latter options (Goldman 2009). And finally, and most importantly, even if the intuitive notion of evidence were tied to the concept of knowledge, it is hard to see how that could be a good thing, given the way that the concept of knowledge arbitrarily privileges causal sources of belief over others.

13 Conclusion

Many recent epistemologists see their task as developing theories that provide 'the best explanation of all of the intuitive data' (Smithies 2012, p. 266). Because of this, they are blind to the possibility that these intuitive data are flawed and our epistemological practices defective.

Reflection on the role of statistical evidence in criminal law makes it clear how intuitions about knowledge can distort our thinking and thwart our enterprises. The criminal courts aim to convict the guilty and free the innocent. Their intuitive demand for knowledge inevitably makes them worse at this than they need be.

The point generalizes. Our traditional attachment to knowledge makes us favour beliefs with certain causal antecedents over others that are just as good. This arbitrarily hampers all our theoretical and practical reasoning. The concept of knowledge needs to be recognized for what it is, a relic of a prehistoric age that serves only to hinder the pursuit of truth.

Acknowledgements I would like to thank the following for help with the ideas in this paper: Yuval Abrams, Maria Alvarez, Maria Rosa Antognazza, Alexander Bird, Bill Brewer, Patrick Butlin, Julien Dutant, David Enoch, Kati Farkas, Clayton Littlejohn, Matteo Mameli, Cosim Sayid, Nick Shea, Matt Soteriou, Levi Spectre, and Scott Sturgeon.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Aikhenvald, A. (2004). *Evidentiality*. Oxford: Oxford University Press.
- Cohen, L. (1977). *The probable and the provable*. Oxford: Clarendon Press.
- Craig, E. (1990). *Knowledge and the state of nature*. Oxford: Clarendon Press.
- DeRose, K. (2002). Assertion, knowledge and context. *Philosophical Review*, 111, 167–203.
- Enoch, D., & Fisher, T. (2015). Sense and “sensitivity”: Epistemic and instrumental approaches to statistical evidence. *Stanford Law Review*, 67, 557–611.
- Enoch, D., Fisher, T., & Spectre, L. (2012). Statistical evidence, sensitivity and the legal value of knowledge. *Philosophy and Public Affairs*, 40, 197–224.
- Gardiner, G. (2018). Legal burdens of proof and statistical evidence. In Coady, D. & Chase, J. (Eds.), *Routledge handbook of applied epistemology* (pp. 179–195). Abingdon: Routledge.
- Greco, J. (2010). *Achieving knowledge: A virtue-theoretic account of epistemic normativity*. Cambridge: Cambridge University Press.
- Goldman, A. (2009). Williamson on knowledge and evidence. In D. Pritchard & P. Greenough (Eds.), *Williamson on knowledge*. Oxford: Oxford University Press.
- Harman, G. (1986). *Change in view: Principles of reasoning*. Cambridge: MIT Press.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Hawthorne, J., Rothschild, D., & Spectre, L. (2016). Belief is Weak. *Philosophical Studies*, 173, 1393–1404.
- Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *Journal of Philosophy*, 105, 571–590.
- Hyman, J. (1999). How knowledge works. *Philosophical Quarterly*, 49, 433–451.
- Hyman, J. (2011). Acting for reasons: Reply to dancy. *Frontiers of Philosophy in China*, 6, 358–368.
- Jeffrey, R. (1965). *The logic of decision*. New York: McGraw-Hill.
- Joyce, J. (2004). Williamson and evidence and knowledge. *Analytic Philosophy*, 45, 296–305.
- Kaplan, M. (2003). Who cares what you know? *Philosophical Quarterly*, 53, 105–116.
- Krupenye, C., Kano, F., Hirata, S., Call, J., & Tomasello, M. (2016). Great apes anticipate that other individuals will act according to false beliefs. *Science*, 354, 110–114.
- Laudan, L. (2006). *Truth, error, and criminal law: An essay in legal epistemology*. Cambridge: Cambridge University Press.
- Lempert, R. (1977). Modeling relevance. *Michigan Law Review*, 75, 1021–1057.
- Littlejohn, C. (2017). Truth, knowledge, and the standard of proof in criminal Law. *Synthese* 1–34.
- Magnus, P., & Cohen, J. (2003). Williamson on knowledge and psychological explanation. *Philosophical Studies*, 116, 37–52.
- Nesson, C. (1979). Reasonable doubt and permissive inferences: The value of complexity. *Harvard Law Review*, 92, 1187–1225.
- Onishi, K., & Baillargeon, R. (2005). Do 15-month-old infants understand false beliefs? *Science*, 308, 255–258.
- Papineau, D. (1999). Normativity and judgement. *Aristotelian Society Supplementary Volume*, 73, 17–43.
- Papineau, D. (2013). There are no norms of belief. In T. Chan (Ed.), *The aim of belief*. Oxford University Press.
- Pritchard, D., Turri, G., & Adam Carter, J. (2018). The value of knowledge. *Stanford Encyclopedia of Philosophy*.
- Russell, B. (1912). On the notion of cause. *Proceedings of the Aristotelian Society*, 13, 1–26.
- Smith, M. (2018). ‘When does evidence suffice for conviction? *Mind*, 127, 1193–1218.
- Simon, R. (1969). ‘Judges’ translations of burdens of proof into statements of probability. In J. Kennelly & J. Chapman (Eds.), *The trial lawyer’s guide*. Newcastle: Callaghan.
- Smithies, D. (2012). The normative role of knowledge. *Nous*, 46, 265–288.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Oxford University Press.
- Steele, K. MS ‘Crime, punishment, and “individualized” evidence’.
- Strawson, P. (1962). Freedom and resentment. *Proceedings of the British Academy*, 48, 1–25.

- Thomson, J. (1986). Liability and individualized evidence. *Law and Contemporary Problems*, 49, 199–219.
- Tomasello, M., & Moll, H. (2015). Why don't apes understand false belief? In M. Banaji & S. Gelman (Eds.), *Navigating the social world*. Oxford: Oxford University Press.
- Tribe, L. (1971). Trial by mathematics: Precision and ritual in the legal process. *Harvard Law Review*, 84, 1329–1393.
- Unger, P. (1975). *Ignorance: A case for skepticism*. Oxford: Oxford University Press.
- Wellman, H., Cross, D., & Watson, J. (2001). Meta-analysis of theory-of-mind development: The truth about false belief. *Child Development*, 72, 655–684.
- Wells, G., & Hasel, L. (2007). Eyewitness identification: Issues in common knowledge and generalization. In E. Borgida & S. Fiske (Eds.), *Beyond common sense: Psychological science in the courtroom*. London: Wiley-Blackwell.
- Williams, B. (1978). *Descartes*. Upper Saddle River: Harvester Press.
- Williamson, T. (1996). Knowing and asserting. *Philosophical Review*, 105, 489–523.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Yates, S. Q. (2017). *Eyewitness identification: Procedures for conducting photo arrays*. Falls Church: Department of Justice.
- Zagzebski, L. (2003). The search for the source of epistemic good. *Metaphilosophy*, 34, 12–28.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.