



Reasonable Doubt from Unconceived Alternatives

Hylke Jellema¹

Received: 23 April 2021 / Accepted: 16 April 2022 / Published online: 24 May 2022
© The Author(s) 2022

Abstract

In criminal trials, judges or jurors have to decide whether the facts described in the indictment are proven beyond a reasonable doubt. However, these decision-makers cannot always imagine every relevant sequence of events—there may be unconceived alternatives. The possibility of unconceived alternatives is an overlooked source of reasonable doubt. I argue that decision-makers should not consider the defendant’s guilt proven if they have good reasons to believe that plausible, unconceived scenarios exist. I explore this thesis through the lens of the two most influential accounts of rational criminal proof—Bayesian and explanation-based approaches. I draw on related ideas from the philosophy of science to show why and when unconceived alternatives lead to reasonable doubt on either account.

1 Introduction

In criminal trials, judges or jurors decide whether the prosecution succeeded in proving the facts described in the indictment. To make this decision, they often compare competing ‘scenarios’ or ‘stories’—i.e. sequences of events such as ‘John stabbed Mary to death during an attempt to rob her’ (Pennington & Hastie, 1993). These scenarios may be formulated by one of the parties, by investigators or even by decision-makers themselves. However, it is not always possible to imagine every possible scenario—there can be *unconceived alternatives*. Most of these alternatives will be not be worth considering—e.g. that it was actually Queen Elizabeth II who killed Mary. However, there may also be one or more plausible, overlooked alternatives. To give an example, someone else could have had the motive, means and opportunity to kill Mary.

Decision-makers often treat the limited number of possible scenarios that they consider as an exhaustive set of possibilities. This presumption of exhaustiveness is, I will argue, necessary. Without it, we cannot rationally come to a conclusion about

✉ Hylke Jellema
h.jellema@rug.nl

¹ Faculty of Law, University of Groningen, Oude Kijk in ’t Jatstraat 26, 9712 EK Groningen, The Netherlands

whether it is probable that the defendant is guilty. However, there can also be cases where this presumption of exhaustiveness is unjustified, because the decision-maker has good reasons to suspect that plausible alternatives exist. The main contention of this article is that, in those cases, reasonable doubt exists about the defendant's guilt. More precisely, a decision-maker should not consider the alleged acts proven if their evidence supports the belief that plausible, unconceived scenarios likely exist in which the defendant did not commit these acts. Conversely, if the decision-maker lacks good reasons for the belief that such alternatives exist, they may presume their set of explanations to be exhaustive.

Parties sometimes invoke the possibility of unconceived alternatives in court as a source of doubt. For instance, Allen & Pardo (2019, 23) cite various scholars who point out that "there are cases where defendants argue along the lines of 'something else happened' without offering any specific alternatives." However, legal evidence scholars have paid little attention to the rationality of such arguments. When and why it is *reasonable* to doubt the defendant's guilt based on the possibility of unconceived alternatives is therefore an open question, which I tackle in this article.

To make my argument as precise as possible, I cast it in the language of the two most popular models of rational criminal proof: the Bayesian and explanation-based accounts, and their associated interpretations of the proof beyond a reasonable doubt standard. An additional benefit of using these models is that they allow me to connect the ideas that I develop here to existing work on unconceived alternatives. Specifically, philosophers of science have thought extensively about unconceived alternatives as a source of doubt about our current scientific theories. They regularly use Bayesian or inference to the best explanation-based accounts of rational scientific inference to develop their arguments. As I will show, many of their ideas can be analogously applied to the legal context.

The next section explains the two frameworks and their associated interpretations of proof beyond a reasonable doubt (Sect. 2). I then turn to why the possibility of unconceived alternatives leads to reasonable doubt on either interpretation. My argument has four steps: First, for both Bayesians and explanationists, the *discovery* of new scenarios may lead to reasonable doubt (Sect. 3). Second, if we have good reasons to suspect that we *could* discover such alternatives, we should take their potential epistemic consequences into account if we want to meet the beyond a reasonable doubt standard. However, third, anticipating the epistemic consequences that such unconceived alternatives are likely to have if we were to discover them is often impossible (Sect. 4). Fourth, we therefore cannot consider guilt proven beyond a reasonable doubt if we are not justified in believing that we have considered all plausible scenarios (Sect. 5). In Sect. 6 I explore when we may justifiably presume that there are no further alternatives. I reject Amaya's (2009) responsibility account and propose an evidentialist answer instead. On my account, we answer the aforementioned question by weighing our evidence regarding the existence of such

alternatives.¹ I distinguish several types of evidence for or against the existence of unconceived alternatives.

2 Reasonable Doubt and Rational Criminal Proof

The term reasonable doubt usually refers to a legal concept used in common law systems, such as the American and the English. These systems use ‘proof of guilt beyond a reasonable doubt’ as their proof standard for criminal trials. One way to approach the problem of unconceived alternatives would therefore be to dissect how specific common law courts interpret the reasonable doubt standard. Such a doctrinal analysis might show that, according to these interpretations, a serious possibility of unconceived alternatives ought to lead to acquittal. However, I do not believe that such an approach would be fruitful. First, it is notoriously unclear how common law courts interpret this standard. The United States Supreme Court has refused to define the term ‘reasonable doubt’ (Strong & Broun, 1992, p. 447). To the extent that other common-law courts *have* defined the standard, they do so in varying and (often deliberately) vague ways. That is why Laudan (2006, pp. 295–6) calls the standard “obscure, incoherent, and muddled”. Furthermore, even in systems that do not have the reasonable doubt standard, such as many European civil law systems, courts should not convict defendants if they have serious doubts about their guilt. When I talk about the doubt arising from unconceived alternatives, my conclusions pertain to those systems too.

In this article I instead use an *epistemic* approach. On an epistemic view of criminal proof the principal goal of legal trials is drawing accurate factual conclusions and, to the extent that factual errors are made, distributing those errors fairly (e.g. Dworkin, 1985; Goldman, 2001; Stein, 2005). With respect to this last point, false convictions are generally considered to be much worse than false acquittals (Epps, 2014). This is captured in Blackstone’s well-known maxim that “it is better that ten guilty persons escape than that one innocent suffer” (Blackstone, 1962).² On the epistemic approach, the key purpose of the reasonable doubt standard is to distribute errors fairly. By setting a high standard for conviction—that the defendant probably committed the alleged acts—the probability of false convictions is diminished at the expense of a higher probability of false acquittals (Hamer, 2004; Laudan, 2006, 58).

I will argue that, if a decision-maker has good grounds to suspect plausible, unconceived alternatives, then they cannot justifiably believe that the defendant is very probably guilty. Conviction is then unwarranted on epistemic grounds. To make this claim

¹ In criminal cases, some evidence, such as DNA evidence, may directly confirm or disconfirm the ‘ultimate hypothesis’ – that the defendant committed the acts alleged in the indictment. Other evidence may inform us about the quality of the first kind of evidence. For instance, an expert witness may testify on the quality of a DNA sample. The evidence for or against the existence of unconceived alternatives is of the second kind, as it relates to the question whether our evidence of the first kind is good enough to come to a conclusion about the ultimate hypothesis.

² Note that this ratio should not be taken too literally; various other ratios have been suggested.

more precise, let us now look at the two most influential epistemic accounts of criminal proof and how they interpret the reasonable doubt standard.

2.1 Bayesianism

Over the past decades, Bayesian approaches to rational criminal proof have become increasingly popular (Fenton et al., 2016). According to such approaches to reason about criminal evidence means to reason probabilistically. More precisely, on the Bayesian account, legal decision makers are taken to have degrees of belief in the propositions under consideration, which they base on the available evidence (Urbaniak & Di Bello, 2021). These degrees of belief can be modeled as a probability between 0 and 1. On the Bayesian view, our degrees of belief are irrational if they are not in line with the precepts of probability theory. One key demand for our degree of belief in any hypothesis to be rational is that, upon receiving evidence, decision-makers should update it in accordance with Bayes' formula:

$$P(H|E) = \frac{P(E|H) * P(H)}{P(E)}$$

The most relevant part of this formula for the discussion ahead is the denominator $P(E)$, also called the *marginal likelihood*. This marginal likelihood can be further decomposed into:

$$P(E) = P(H_1)P(E|H_1) + \dots + P(H_n)P(E|H_n) + P(H_c)P(E|H_c)$$

Here $H_1 \dots H_n$ are our conceived hypotheses and H_c is the *catch-all hypothesis* which expresses that 'none of our conceived hypotheses is true'. The catch-all therefore conveys the possibility of further, unconceived alternatives. So, in order to determine the value of $P(E)$ we need to determine the likelihood and prior probability of this catch-all. I will come back to why this is problematic in Sect. 4.

I should mention that, in legal settings, Bayesians typically use the 'odds' version of Bayes' formula:

$$\frac{P(H_1|E)}{P(H_2|E)} = \frac{P(E|H_1) \times P(H_1)}{P(E|H_2) \times P(H_2)}$$

This equation expresses the *relative* posterior probability of two hypotheses—i.e. how much more probable is H_1 than H_2 , conditional on the evidence. It does not contain a catch-all hypothesis. However, this relative probability can *only* be converted to an absolute posterior probability of the two hypotheses if H_1 and H_2 are mutually exclusive and exhaustive. We will end up denying this presumption if we suspect unconceived alternatives.

2.2 Bayesian Reasonable Doubt

On the Bayesian account, proof beyond a reasonable doubt requires (at least) that the posterior probability of guilt to meet some threshold (Gardiner, 2019). To put it differently, the probability of the hypothesis that the defendant is guilty should be sufficiently high conditional on the evidence in the case.³ How high should it be? Various authors have proposed values such as 0.9, 0.95 or even 0.99 (Hamer, 2004). Which value is desirable depends on the negative utilities that our legal system assigns to different types of errors (see the remark above on fair error distribution). What matters is that, for proof beyond a reasonable doubt, the decision-maker ought to be highly confident that the defendant committed the alleged acts and that this confidence reflects the available evidence.

2.3 Explanation-Based Approaches

Explanation-based approaches frame criminal evidence in terms of competing explanations of the evidence. On this view, decision makers in criminal cases do or should reason abductively, i.e. ask how well each of the competing scenarios explains the evidence in the case and then accept one (or none) as true.

Various authors advocate explanation-based approaches, including Josephson (2001) who suggests an abductive theory of criminal proof; Pardo and Allen (2008), Allen and Pardo (2019) *Relative Plausibility Theory*; Pennington & Hastie (1993) who give a descriptive, psychological approach—the *Story Model*; van Koppen (2011) who develops a normative version of the Story Model; Bex (2011) who combines the story model with an argumentation-based approach; and Amaya (2009, 2015) who suggests a coherence-based theory. I do not want to delve too deeply into the distinctions between these various proposals. What they share is the commitment that evidential reasoning in criminal cases is reasoning about competing, causal explanations of the evidence. Or, as Pardo and Allen (2008, pp. 229–33) put it: “explanation as a guide to inference.” In criminal cases these explanations typically take the form of ‘scenarios’ or ‘stories’, which are essentially time lines—series of causally connected events that explain what happened in a case. Empirical research suggests that judges and juries commonly construct such stories and compare them to see which they find most believable when deciding on a case (Pennington & Hastie, 1993).

On explanation-based accounts, we evaluate these scenarios through *explanatory reasoning*. This means asking how well each scenario explains the evidence. In a broad sense, a scenario explains the evidence well to the extent that, if we imagine the scenario to be true, the evidence in the case makes sense to us. Various authors have articulated further criteria for determining an explanation’s ‘explanatory quality’ or ‘plausibility’. To give an example, Pardo and Allen (2008, p. 230) propose

³ Some argue that this interpretation is problematic, e.g. because it leads to counterintuitive outcomes in cases of naked statistical evidence or that it is too subjective (Ho, 2015). I will leave these considerations aside.

that “all other things being equal [an explanation is] better to the extent that it is consistent, simpler, explains more and different types of facts (consilience), better accords with background beliefs (coherence), is less ad hoc, and so on; and is worse to the extent that it betrays these criteria.” While such an attempt to spell out explanatory quality leaves much to be desired in terms of clarity and specificity, they give some indication on what makes an explanation ‘good’ or ‘plausible’.

2.4 Explanationist Reasonable Doubt

Explanationists have not engaged extensively with the question how we should interpret the proof beyond a reasonable doubt standard. To the extent that they do, they typically adopt the following interpretation from Pardo and Allen (2008, pp. 238–9):

In criminal cases [decision-makers should infer] the defendant’s innocence whenever there is a sufficiently plausible explanation of the evidence consistent with innocence (and ought to convict when there is no plausible explanation consistent with innocence assuming there is a plausible explanation consistent with guilt).⁴

In other words, when we determine whether the guilt of the defendant has been proven, we should check whether there is a believable case for guilt and no believable case for innocence.⁵ The former is typically presented by the prosecution, while the defense often provides the latter.

Though the explanationist approach is distinct from the Bayesian, it shares the commitment that a defendant can *only* be convicted if the probability of their guilt is sufficiently high. As Allen and Pardo (2017, p. 1580) put it: “The explanatory account shares the same ends or goals as probabilistic approaches, which have to do with various policy judgments about the likelihood of disputed facts and allocating the risk of error between the parties.” In other words, on both accounts a high probability of the defendant’s guilt is a requirement for proof of guilt beyond a reasonable doubt for the sake of error distribution.

How the Bayesian notion of probability relates to that of plausibility is a topic of ongoing discussion (see e.g. Biedermann & Vuille, 2018; Allen & Pardo, 2019; Mackor et al., 2021).⁶ However, for the sake of the following discussion I will use these terms as follows: A plausible scenario explains the relevant facts well. For instance, a scenario may be plausible because it offers a coherent, detailed explanation of the most important facts in a case. In contrast, probability is a quantification

⁴ Some of the authors who use this interpretation of the standard are Josephson (2001, p. 1642), Allen (2010), Bex & Walton (2012, p.120).

⁵ Pardo and Allen do not specify when an explanation is *sufficiently* plausible. I leave this issue aside.

⁶ The relationship between probability and explanatory quality is also an ongoing discussion in the philosophy of science. For instance, Lipton (2003) argues that, that though explanatory quality (which he calls ‘loveliness’) can be a guide to the probability of an explanation (its ‘likeliness’), the two should not be equated. In contrast, other philosophers of science suggest that explanatory goodness may be translated into a probabilistic notion, and that that explanation-based approaches can therefore be fitted into the Bayesian framework see e.g. Douven (2017).

of how confident an agent is that some hypothesis—for instance a scenario—is true. The two notions are closely related because—all other things being equal—the more plausible a scenario is, the more probable we should consider it. However, what is important for the discussion ahead is that a plausible scenario is not *always* probable. For instance, as I will discuss later, a plausible scenario can be improbable if it has many plausible competitors, while a somewhat implausible scenario may occasionally be quite probable if we are certain that there are no better alternatives.

3 Bad lots and Reasonable Doubt from New Alternatives

In this section, I begin sketching the problem of unconceived alternatives. Few legal scholars have paid attention to this problem. Those who do discuss it, take it to be a more narrow problem than it actually is. Specifically, they frame unconceived alternatives as being problematic only for *inference to the best explanation* (IBE)-based models of criminal proof—a subset of explanation-based approaches. This is called the *bad lot problem* (Amaya, 2009; Jellema, 2020; Laudan, 2007; Ribeiro, 2018).⁷ The bad lot objection invokes the possibility of unconceived alternatives to argue that IBE does not reliably lead to accurate outcomes. According to Amaya (2009, p. 154) it is “the most serious problem that a model of IBE for law has to face.” However, while I agree that the bad lot problem is a major objection to IBE-based approaches, unconceived alternatives are problematic for other major accounts of rational criminal proof as well. It is this broader problem that I want to discuss here.

Let me begin by explaining the bad lot objection against IBE-based models. Inference to the best explanation is a specific approach to explanatory inference, where we infer that the explanation which best explains the evidence is (probably) the true explanation (Lipton, 2003). The bad lot objection begins with the observation that IBE uses comparative reasoning—that a hypothesis is the best available—to arrive at an absolute verdict—that this hypothesis is true (Douven, 2017, p. 9). However, such a conclusion is only warranted if we may presume that the true hypothesis is probably among those considered.⁸ If we drop this presumption, we end up with the merely comparative conclusion that one explanation is better than those we have come up with so far. Or, in terms of criminal cases, without this presumption, IBE only justifies the conclusion that one scenario is the best out of those that have been presented in court. Any such comparative conclusion is obviously insufficient to ground conviction. For instance, it might even mean basing a conviction on a poor scenario that is probably untrue, simply because this scenario is the best among the lot.

⁷ This problem was originally formulated by philosophers of science (Sklar, 1981; Van Fraassen, 1989, p. 142–3).

⁸ Or, as Van Fraassen (1989, p. 143) puts it: “[T]o believe the best explanation (...) requires a step beyond the comparative judgment that the hypothesis is better than its actual rivals. (...) For me to take it that the best of set X be more likely than not, requires a prior belief that the truth is already more likely to be found in X, than not.”

Though the bad lot problem poses a difficulty for IBE-based models of criminal proof, it does not affect the Bayesian, nor many of the explanation-based models discussed in the previous section. After all, these models do not rely on comparative reasoning in the same way that IBE does. On these models we do not go from the conclusion that one explanation or hypothesis is the best available to the conclusion that it is probably true.⁹ For instance, recall that on many explanation-based accounts, guilt is proven beyond a reasonable doubt if *there is a plausible guilt explanation and no plausible innocence explanation*. Yet as Pardo & Allen (2007)—the originators of this interpretation—note in response to the bad lot objection, this means that we are not choosing the best explanation of the available ones. Rather, we check whether the prosecution has provided a scenario that is sufficiently plausible to ground a conviction and whether the defense (or anyone else)¹⁰ has produced a scenario that is sufficiently plausible to create a reasonable doubt.

Still, unconceived alternatives also pose a problem for Allen and Pardo's interpretation of the reasonable doubt standard, as well as for the Bayesian account. Specifically, within these frameworks, evidential inference also requires the presumption that we have not overlooked plausible alternative scenarios. If this presumption is unjustified (because we have good reasons to believe that plausible unconceived alternatives exist) guilt cannot be proven beyond a reasonable doubt. This is the broader worry of unconceived alternatives.

My argument that this presumption is required begins with the observation that the discovery of *previously* unconceived alternatives can create a reasonable doubt on both the explanation-based and Bayesian account. To start with the former, consider the two demands of proof beyond a reasonable doubt on this interpretation: (i) There must be a plausible explanation that implies the defendant's guilt, and (ii) there must not be a plausible explanation that implies his innocence. Suppose that, at some point, both these conditions are met. A subsequent discovery of a previously unconceived alternative could result in either (or both) of these conditions no longer being satisfied. This is most obvious with the second condition: the new alternative might both be plausible and imply the defendant's innocence. But even the first condition might be undermined. We may deem the prosecution's scenario plausible *because* there are no plausible alternatives. Criminal cases often concern unusual situations, where events happen that we might not otherwise regard plausible. For instance, people may do strange things when acting impulsively under stress. A story containing such actions could—at first sight—not seem to explain the available evidence well. But if there are no realistic alternatives, this might us consider the story believable after all. This also means that if we discover that there *are* one or more plausible rivals, this can lead us to considering a scenario implausible.

⁹ This is not to say that such comparative reasoning never plays a role on these models. As I will argue below, how one hypothesis fares with respect to its competitors often *is* relevant for the question how probable it is. Nonetheless, this comparison is not an *inherent* part of these models, and hence these models avoid the narrower bad lot objection, which targets the move from a relative to an absolute conclusion, as discussed above.

¹⁰ This is similar to the way that Lipton (2003) suggests that IBE can overcome the bad lot problem, namely by adding the demand that the best explanation should also be 'good enough'.

So, on explanation-based accounts, the discovery of a previously unconceived scenario may overturn the case for guilt. What about the Bayesian approach? At first sight, this approach may seem to avoid the problem of newly discovered alternatives altogether. After all, on the Bayesian account when we consider the probability of some hypothesis *H*, we typically take into account both that hypothesis and its negation. For instance, in their review article on Bayesian inference in legal settings Fenton et al., (2016, p. 53) begin by introducing the basics of the Bayesian approach. They write: “A hypothesis is a statement (...) whose truth value we seek to determine. Examples include: a. ‘Defendant is innocent of the crime charged.’ (...) b. ‘Defendant was the source of DNA found at the crime scene.’ (...) The alternative hypothesis is a statement that is the negation of a hypothesis.” When we cut up the hypothesis space this way, we seemingly avoid the problem of unconceived alternatives. After all, either the defendant is innocent or she is not. Either she is the source of the DNA or she is not. There are no other options to consider as these hypotheses jointly exhaust the hypothesis space.

The problem for the Bayesian is that the above picture is unrealistic. Humans typically do not reason in terms of such general, exhaustive hypotheses, nor may we expect them to. For instance, legal decision-makers typically consider a small set of specific scenarios which helps them make sense of the available evidence (Pennington & Hastie, 1993). One key reason for this is that to adequately consider such mutually exclusive and exhaustive hypotheses would require the fact-finder to consider all of the countless (or even infinite) possibilities in the hypothesis space. This will often be an impossible feat for even the smartest person. Indeed, contemporary Bayesians also frequently note that thinking in terms of mutually exclusive and exhaustive hypotheses is not always realistic. For instance, Urbaniak and Di Bello (2021) point out that considering the negation of a specific hypothesis (e.g. ‘the defendant did not hit the victim in the head’) “can be unhelpful in assessing the evidence.” They suggest that determining the probability of such a negation will involve on drawing up a more specific scenario, which includes details such as whether the defendant was at the crime scene. Several authors have recently developed Bayesian frameworks that reflect this point, by incorporating the notion of scenario-based reasoning (e.g. Cheng, 2013; Dahlman, 2019; Urbaniak, 2018; Vlek et al., 2013).

If we consider only a limited set of scenarios, then the possibility of unconceived alternatives rears its head. Hence, the Bayesian also faces the possibility of discovering new, previously unconsidered alternatives. Furthermore, just as on explanation-based accounts, the discovery of such alternatives may lead to reasonable doubt for Bayesians. Firstly, it is a generally accepted fact that discovering that there are further, previously unconsidered alternatives, may change the probabilities that we assign to the other, already conceived hypotheses. The question how Bayesians should account for this is called ‘the problem of new theories’ (Talbot, 2008).¹¹ Secondly, evidence that there are *no* plausible alternatives can sometimes raise the

¹¹ This is related to, but not identical to, the well-known problem of ‘old evidence’. The old evidence problem refers to the question how evidence that has been known for some time can provide support for existing theories, when we discover that there is a logical relation between the two (see Talbot, 2008, 6.2 for further discussion).

probability of a hypothesis (Dawid et al., 2015).¹² Unconceived alternatives therefore pose the same problem for Bayesians as for explanationists: their discovery may overturn the case for guilt (and sometimes discovering that there are no alternatives may be the reason *why* guilt is proven beyond a reasonable doubt).

The *discovery* of a new alternative may therefore lead to reasonable doubt on either account. Yet the central thesis of this article is stronger than this. In some situations, the fact that we *could* realistically discover a plausible alternative should be a cause of reasonable doubt. Of course we can never be certain that we have exhausted all possibilities, even in the most clear-cut of cases. The situations that I am concerned with are those where we have *good reasons* to suspect that there are unconceived alternatives—i.e. that there is some relevant possibility that we have overlooked. A defense lawyer could reasonably invoke such grounds to argue that guilt has not been proven beyond a reasonable doubt. I will develop a fuller account of what such reasons are in Sect. 6. For now, let me just mention two examples. First, we may reasonably suspect that there are further unconceived alternatives if we know that the police focused only on a single suspect throughout their investigation, neglecting to look for evidence that there could be others. Second, all our currently conceived scenarios might be implausible. For instance, we might discover that both the defense and prosecution scenario contain illogical time leaps. In that case we may also justifiably presume that there has to be some further alternative that we overlooked. This could either be a completely new scenario, or a variation on an old scenario that explains the difficulty away (e.g. a scenario which posits that the evidence which currently makes our scenario implausible is misleading).

In the next two sections I argue that, when we have good reasons to suspect plausible unconceived alternatives, we are not justified in considering guilt proven to a high degree of probability. We may then, at best, be justified in drawing the conclusion that some scenario fares *comparatively well* in relation to the others that we have conceived of.¹³ The reasoning for this is as follows: because discovering plausible alternatives may lead to reasonable doubt, we should include this possible impact in our inferences about whether guilt has been proven beyond a reasonable doubt if we believe that unconceived alternatives probably exist. However, as I will argue in the next section, including this potential impact is often impossible.

4 The Difficulty with Considering the Unconceived

Imagine that we have good reasons to suspect that there are plausible unconceived alternatives. Furthermore, suppose that we want to take this fact seriously and take it into account when deciding whether the defendant's guilt has been proven beyond a reasonable doubt. The answer is, at first sight, straightforward. Explanation-based

¹² In fact, the earlier Bayesian approach on which decision-makers reason about general, exhaustive hypotheses is arguably unrealistic partially because it does not leave room for the epistemic consequences of unconceived alternatives.

¹³ Which is precisely the charge that critics have leveled against IBE by means of the bad lot problem.

accounts make inferences in terms of explanations and their degree of plausibility. Bayesian inference concerns hypotheses and their prior probabilities and likelihoods. So, on either account we could treat the possibility that there are further alternatives as an explanation or hypothesis and we assign it a degree of plausibility or a prior probability and likelihood. For instance, on the explanation-based account we might then end up with the set of scenarios {‘John killed Mary’, ‘William killed Mary’, ‘Something else happened’}. This last scenario expresses the possibility of unconceived alternatives. We might then assign each scenario a degree of plausibility. If the last scenario turns out to be plausible, we could treat it as an innocence scenario and acquit both John and William on the grounds that there is a plausible alternative that does not imply their guilt.

Sadly, this proposal does not work. The reason for this is that the possibility of unconceived alternatives cannot be evaluated in the same way that we usually evaluate scenarios. To see why, consider an example due to Dellsén (2017a, p. 37):

Suppose you come home one day to find the front door open and the lock broken. Furniture is overturned, the contents of the shelves are on the floor, and valuables are missing. One explanation is that someone broke in and stole your belongings, making a mess in the hurried process.

Call the scenario mentioned in the final sentence the break-in hypothesis. Now consider the negation of this, the no-break-in hypothesis—that it is not the case that someone broke in, stole your belongings, and made a mess. Or, to put it differently, ‘something else happened’. On the explanationist approach, we should evaluate these explanations in terms of how well they explain the evidence. This is unproblematic for the break-in hypothesis; the explanation that someone broke into your house explains various facts such as why the lock is broken and why the furniture is overturned. The no-break-in hypothesis, in contrast, offers no explanation of these facts. It offers us no understanding of why the lock was broken or why the furniture was overturned.

As Dellsén (2017a, p. 33; 2018, pp. 1758–9) points out, if we engage in inference to the best explanation, we should therefore trivially favor the break-in hypothesis (or really any explanatory hypothesis) over its negation when we ask which explanation is the best.¹⁴ This even goes for poor explanations. For example, consider the scenario that ‘aliens dancing wildly to ska music turned over the furniture and broke the lock’. If we evaluate explanations in terms of their explanatory quality we should favor the alien scenario over its negation. After all, the latter has no explanatory quality whatsoever—it offers no account of how the mess in the house was created—whereas the former does (however bad an explanation it may be). But if even a bad explanation ranks above the no break in-hypothesis, then no matter how probable it actually is in this case that ‘something else happened’, this possibility will never be accepted as the best explanatory hypothesis (assuming we have formulated

¹⁴ Dellsén makes this point when objecting to Lipton’s (2003) response to the bad lot problem. Lipton’s response is more complex than I have room to explain here, but it ultimately fails because of this point—that we cannot properly compare explanatory hypotheses to negations ().

at least a single possible explanation). This would, for instance mean that if the alien hypothesis is our *only* conceived explanation, we would then have to accept it as true on IBE.¹⁵ But, obviously, this does not mean that the alien hypothesis is probably true.

What about explanation-based accounts that do not rely on IBE? In the last section I pointed out that not every explanation-based approach suggests that we should choose the best explanation. Nonetheless, the above point is equally problematic for those explanation-based accounts which state that proof beyond a reasonable doubt requires a plausible guilt explanation and no plausible innocence explanation. No matter how probable it is that we have missed something, the hypothesis that ‘there are unconceived alternatives’ will never be counted as a plausible alternative scenario if we use explanatory reasoning. Unconceived alternatives simply do not show up in regular explanatory inference. For that to happen, we would have to make it more specific *what* might have happened, if not our conceived scenarios. And in doing so, we would simply be coming up with another alternative explanation—we would not be assigning a degree of plausibility to the unconceived. We cannot compare the unconceived to the conceived precisely because we know so little about it.

At this point you might suspect that the problem is with explanatory reasoning and that the Bayesian account will therefore fare better. However, as various authors have noted over the years, Bayesianism is ill-equipped to deal with the possibility of unconceived alternatives. To see this, remember how Bayesians may include the possibility that there is some unconceived alternative by including a so-called ‘catch-all hypothesis’ which expresses the hypothesis that ‘none of our current hypotheses is true’. As mentioned earlier, in the standard formulation of Bayes’ formula, the catch-all hypothesis is included in the denominator, $P(E)$. In the odds-version, we can only draw conclusions about the absolute probabilities of the hypotheses under consideration if these hypotheses are mutually exclusive and exhaustive. To achieve this, we can include the catch-all, which would be in the denominator, where assigning a likelihood, $P(E|\neg H)$ and prior probability $P(\neg H)$ requires assigning a prior probability and likelihood to the catch-all too. For both versions of the formula, because the catch-all is in the denominator, we can *only* consider H highly probable if the value of the catch-all is low.

In order to know whether H is highly probable or not, we therefore need to determine the likelihood and prior of the catch-all. But here the Bayesian account faces a difficulty. The catch-all is made up of the disjunction of all alternative hypotheses to those in our set of conceived hypotheses $\{H_1 \dots H_n\}$. So, to assign a value to it we need to determine (or at least approximate) the likelihoods and prior probabilities of these alternatives, without knowing any specifics about them.

For some contexts, defining these values for the catch-all is unproblematic. An example of this might be a card game. In that setting we can accurately determine the probability of the hypothesis ‘the next card will *not* be a three of hearts or a four of clovers.’ In criminal contexts we might similarly be able to determine the

¹⁵ Contemporary defenders of IBE have modified their accounts to avoid such a conclusion. For instance, Lipton (2003) adds the demand that the best explanation should also be ‘good enough’.

random-match probability of some DNA sample. However, on the level of scenarios, such values will typically not be so easy to assign. For any given scenario there will be an almost infinite number of alternative scenarios of which we do not know what they look like (after all, they are unconceived). It is difficult to think of a feasible method of determining the likelihood and prior probability of such a set. We simply do not know enough about what the unconceived looks like to assign it a reasonable likelihood and prior probability. As Steele and Stefánsson (2019) put it:

[I]n order for an agent to make sense of a catch-all, she would presumably need to entertain some universal set of possibilities relative to which the catch-all can be defined as the complement of those possibilities she can properly articulate. But it is hard to see how the agent could have access to this universal set of possibilities (...), given that, by assumption, some of these possibilities cannot be articulated.¹⁶

For instance, consider the earlier example of the possible break-in again. How should we assign a likelihood of the evidence given the catch-all in that situation? For the break-in hypothesis we could reasonably say that the likelihood is high because a break-in explains the facts well. What about the no-break-in hypothesis? How can we determine how expected the evidence is given the fact that there was no break in, if we don't know any details about what this no-break-in-situation actually looked like? It seems that there is no meaningful way of doing so. As Salmon (1990, p. 329) writes: "What is the likelihood of any given piece of evidence with respect to the catch-all? This question strikes me as utterly intractable". The same goes for the prior probability of such a catch-all hypothesis. As Sklar (1981, p. 19) writes: "[W]e must distribute a priori probabilities over all the alternative hypotheses to be considered. If there is only a finite set of hypotheses we have in mind, this is easy to do [...]. But if we must keep in mind the infinite and indeterminate class of all possible hypotheses, known and unknown, how can we even begin to assign a priori probabilities to those few hypotheses [...] we do have in mind [...]?"

That it is difficult to assign any definite prior probability to unknown hypotheses becomes especially evident once we consider what happens when we *do* discover such a previously unconceived hypothesis. This is a well-known problem for Bayesians, because it seems that the probabilities of our existing hypotheses should change, but there is no new evidence to conditionalize on (Talbot, 2008, Sect. 6.2). The most common way in which Bayesians deal with this is by 'shaving off' (Earman, 1992; Wenmackers & Romeijn, 2016). Simply put, a newly conceived hypothesis 'steals' its new probability from the catch-all; the latter's probability is lowered so that we can assign a positive probability to the new hypothesis. Suppose that we originally had the set of hypotheses $H_1 \dots H_n$ and the catch-all H_c . Now suppose that we conceive of the alternative hypothesis H_{n+1} . We have to assign this hypothesis some non-zero probability. Because the set $\{H_1 \dots H_n, H_c\}$ was mutually exclusive

¹⁶ Similarly, Bradley (2017, p. 255) writes: "[G]iven that we don't know anything about the prospects that we are potentially unaware of, on what basis are we to determine (...) what probability we should assign to the catch-all prospect?"

and exhaustive, their probabilities by definition added up to 1. To assign the new hypothesis, H_{n+1} a non-zero probability therefore requires that we lower the probability of the original set. We may do so by assigning the new catch-all, H_c , a lower probability, so that the probabilities of the resulting set add up to 1.¹⁷ However, as Earman (1992, p. 195) argues, this means that as more and more new hypotheses are discovered and the catch-all becomes smaller, we can only assign them smaller and smaller probabilities when ultimately the point may be reached “where the new theory has such a low initial probability as to stand not much of a fighting chance.” So, if we want to leave open the possibility of discovering many plausible alternatives, the prior probability of the catch-all should be very high. Yet this means that the prior probability of our *conceived* hypotheses has to be very low, regardless of how plausible these hypotheses actually are.

Admittedly, the specific point raised by Earman is often more readily applicable to the scientific context than to the legal. Throughout scientific history, there have been numerous instances of novel theories leading to revisions of the probability of the old theories. Some have argued that we therefore have good reasons to continue to expect further plausible theories to be discovered in the future and that it is therefore rational to assign a low probability to our current best theories too (Stanford, 2006). In contrast, in a specific criminal trial, we typically have little reason to expect numerous plausible alternatives to be discovered.¹⁸ Nevertheless, the question remains how high the prior of the catch-all should be if we want to leave room for the discovery of plausible alternatives, without knowing exactly how plausible these alternatives are. If we make the prior of the catch-all too low, we may not leave enough room to assign newly discovered, plausible alternatives a high enough probability. If we make the prior too high, it becomes difficult to prove anything, even if our current scenarios are highly plausible. Hence, it seems that there is no reasonable way to assign a prior probability to the catch-all.

To clarify these points, consider the analogous problem of *missing evidence*. For instance, imagine that investigators neglect to interview an important witness. Various authors suggest that, if there is a great deal of missing evidence, the case for guilt may lack the appropriate ‘weight’ or ‘evidential completeness’ (Ho, 2015, Sect. 3.3). But how should we account for the weight of the evidence? One idea is that worries about weight can be reflected in how probable the evidence makes the ultimate hypothesis—that the defendant committed the alleged acts. We then take the missing evidence into account by estimating what the influence on the relevant probability could be if we did have it. However, as Nance (2008, pp. 633–9) points out, the problem with this is that we do not know what the content of the missing information is. We may sometimes guess this content, e.g. if we know that one party

¹⁷ Note that by doing so we create a new catch-all which differs from the former as it no longer includes the previously unconceived H_{n+1} .

¹⁸ Though this will, of course, depend on the particulars on the case. For instance, in some cases it may be relatively clear how the crime happened and which people might have committed it—for instance in so-called island-cases, where the number of possible culprits is clearly delineated (see Fenton et al., 2019). In other cases, we may expect there to be many—or even infinite—plausible alternatives which vary with respect to the perpetrator, means, motive, location, time and nature of the crime.

repressed evidence, then that evidence likely supported the other side's case. But, as Nance (2008, p. 274) writes "that gives little hint of an answer to the question, 'By how much?'" Similarly, we cannot reasonably draw inferences about whether unconceived alternatives would support the defendant's guilt or innocence and by how much. We therefore cannot take the impact of these unconceived alternatives into account before we discover them by including a catch-all hypothesis or by adding the scenario that 'something else happened'. There is simply no way to evaluate the unconceived alternatives in a way that would reflect their potential impact. Any value that we assign to the catch-all does not necessarily correspond to what we would assign to these alternatives were we to discover them.

Admittedly, there are situations in which we may justifiably exclude the catch-all from our inferences, even if our set of alternatives is not exhaustive in a strict sense of the word. For example, Fitelson and Thomason, (2008 p. 26) point out that we may sometimes believe that every unconsidered possibility may be so implausible as to be negligible. We can then presume that the prior probability of the unknown hypotheses is (very) low and that the sum of the prior probabilities of the known theories is therefore (very) high. We are then justified in treating our current set of hypotheses as exhaustive. Indeed, this is, I believe, a typical assumption in criminal law, where decision-makers only look at a select few scenarios, but treat those as if they exhaust the hypothesis space. But the question I want to ask is, what if we *cannot* justifiably make this presumption—because we have good reasons to suspect *plausible* unconceived alternatives? In the next section, I argue that we should then have a reasonable doubt about the defendant's guilt.

5 Reasonable Doubt from Unconceived Alternatives

The upshot from the previous section is that we cannot meaningfully take the possibility of unconceived alternatives into account when we assign plausibility or probability to the available scenarios. Yet suppose that we (a) suspect the possible existence of unconceived, plausible alternatives, but (b) do not take this possibility into account when evaluating our set of scenarios. In other words, we let go of the presumption that our set of scenarios is exhaustive. If we do so, we can no longer draw conclusions about the absolute probability of the defendant's guilt, which is what the beyond a reasonable doubt standard requires.

For instance, some propose Bayesian accounts that do not rely on an exhaustive set of hypotheses. Salmon (1990) is an example. He suggests that we should only consider conceived alternatives when evaluating the confirmation of a given hypothesis. But, as Rowbottom, (2019, p. 3) points out in response, if we only consider conceived alternatives, this means letting go of the assumption that we are evaluating whether a theory is "truth-like." Similarly, Wenmackers and Romeijn (2016) propose an 'open minded' version of Bayesianism, which drops the assumption "implicit in standard Bayesianism—that the correct empirical hypothesis is among the ones currently under consideration". However, they admit that their approach "fails to provide us with the required normative guidance" about the absolute confirmation of scientific theories, because it only tells an agent what to believe if she

supposes “without committing to it, that the true theory is among those currently under consideration” (Wenmackers & Romeijn, 2016, p. 1243) In other words, if we drop the presumption of exhaustiveness, we cannot consider the absolute probability of any hypothesis to be high.¹⁹ The demand for proof beyond a reasonable doubt then cannot be met.

What about explanation-based accounts? As mentioned in Sect. 3, the bad lot problem arises for inference to the best explanation when we cannot presume that the true hypothesis is probably amongst those considered—which is precisely the presumption that we are rejecting here.²⁰ However, as I also discussed in that section, many explanation-based accounts do not rely on inference to the best explanation. They require that there is a plausible guilt explanation and no plausible innocence explanation. Nonetheless, without the presumption of exhaustiveness, these accounts also lead only to a kind of comparative conclusions. We then cannot conclude that there are no plausible alternatives. At best we can infer that no plausible alternatives have been presented in court. But we do not care only about what has been presented in court. We also want to know whether what has been presented in court reflects the actual strength of the case for innocence and guilt.

The point of the beyond a reasonable doubt standard—both on the Bayesian and explanation-based account—is to allow conviction only when the defendant probably committed the alleged acts. But we cannot draw such absolute conclusions while at the same time taking the possibility of unconceived alternatives into account. Some may bite this bullet and suggest adopting a comparative account of rational criminal proof.²¹ As far as I am aware, only Cheng (2013) has made such a proposal. He argues that the reasonable doubt standard could be reconceptualized as a Bayesian likelihood-ratio. On his account decision-makers should compare “a single defense narrative of innocence versus a single prosecution narrative of guilt” and consider the latter proven if it has a sufficiently higher likelihood ratio. However, I am not certain that Cheng indeed intends his account to be merely comparative or whether he sees some connection to the absolute probability of guilt. After all, the price of bullet-biting is high, as it severs the connection between the proof standard and its epistemic aim of fair error distribution.²²

So, how do we deal with unconceived alternatives without biting this bullet? The solution that several philosophers propose—and that I also endorse for the legal context—is to treat the exhaustiveness of our set of hypotheses as a prerequisite for our

¹⁹ Nor can we always say that it is low.

²⁰ In response, some philosophers of science have also offered strictly comparative accounts of inference to the best explanation – where we merely infer that one explanation is more likely to be true than the other available ones (e.g. Kuipers, 2004).

²¹ While Allen & Pardo (2019) call their approach the ‘relative plausibility theory’, their account does not lead to relative conclusions in criminal cases. See Sect. 2.2.

²² Outside of the problem of not having mutually exclusive and exhaustive hypotheses, the likelihood ratio also does not take into account the prior odds of the hypotheses. Prior odds are necessary to arrive at posterior odds (Meester & Sjerps, 2004). Cheng (2013) argues that such prior odds do not need to be taken into account in the proof standard. I do not discuss his argument, as it is not germane to the matter at hand.

Bayesian or explanationist inferences to be justified. For example, as Amaya puts it, in order for IBE to work, our set of explanations should be ‘good enough’. For Bayesianism, the story is the same—the presumption that our set of conceived possibilities is exhaustive can be seen as a prerequisite for Bayesian inference (see e.g. Gillies, 2001).

Note that by prerequisite, I do not mean that we first have to determine whether there are unconceived alternatives before we can engage in Bayesian or explanation-based reasoning. As I will argue in the next section, how plausible our current scenarios are also informs us about whether there are unconceived alternatives. Additionally, as discussed earlier, we should sometimes consider a scenario plausible because we are confident that there are no alternatives. So, while these are two separate questions, they are interrelated and may have to be answered in conjunction with one another. The question of unconceived alternatives therefore does not come prior to whether our conceived scenarios are plausible or probable. Rather, what I mean by prerequisite is that *if* we cannot justifiably believe that our set of explanations is sufficiently exhaustive *then* we are also not justified in accepting the conclusions of our Bayesian or explanationist inferences.

6 Justified Belief in No Unconceived Alternatives

Given our human, cognitive limitations, we usually cannot be certain that we have considered every plausible alternative. Nonetheless, we may be *justified* in presuming that we have. The challenge is then to spell out when this presumption is justified. On my account this depends on how strong our *evidence* is for the existence of plausible unconceived alternatives. Before I delve further into what counts as evidence for unconceived alternatives, I first contrast my evidentialist account with Amaya’s (2009) responsibilist framework. She is, to my knowledge, the only other author who has tackled the question how we may reason about unconceived alternatives in the context of legal proof.

6.1 Responsibilism

Amaya discusses the problem of unconceived alternatives when she defends her inference to the best explanation-based account of rational criminal proof from the bad lot problem. In order to avoid bad lots, she proposes that we are only justified in inferring to the best explanation if our set of explanations is ‘good enough’. Without this, she suggests, we may not conclude that the best explanation is likely to be true. Her aim is therefore the same as what I seek to achieve in this article, though her solution differs from mine.

Amaya distinguishes between *responsibilist* and *non-responsibilist* (or *evidential*) views of justification. On the first, justification is about what an agent has done (or failed to do) to ensure that her beliefs are true. On the second, justification is analyzed purely in terms of evidential support. According to Amaya, we ought to be responsibilists with respect to whether our set of explanations is ‘good enough’. This

means “complying with some epistemic duties and exercising a number of epistemic virtues in the course of inquiry and deliberation about factual problems in law” (Amaya, 2009, p. 155).²³ She lists some of these virtues: “open-mindedness in collecting and appraising evidence, perseverance in following a line of inquiry, or readiness to change one’s views in the face of new conflicting evidence” (Amaya, 2009, p. 155). So, whether we may presume that the set of scenarios is good enough depends on whether investigators and prosecutors have acted in an epistemically virtuous way (and nothing else).²⁴

Though Amaya does not argue *why* we ought to be responsibilists when discussing the bad lot problem, she does offer reasons for this elsewhere: while our evidence may support a belief, this may only be the case because this evidence was gathered in an epistemically irresponsible (e.g. biased, lazy) manner (Amaya, 2008). Such a belief is then not actually justified, even though the evidence supports it. We can imagine a similar justification with respect to unconceived alternatives: a scenario that implies defendant’s guilt may be quite probable given our evidence, but only because we tried insufficiently hard to coming up with alternative scenarios. Additionally, a commitment to responsibilism is generally understandable in the context of criminal law. Epistemic duties on the part of the prosecution are crucial in criminal proof. In particular, the prosecution holds the burden of proving the defendant’s guilt. If a prosecutor or investigator has acted epistemically irresponsible, this should undermine the case that they try to build against the defendant.

However, despite this intuitive plausibility, the responsibilist account is unsatisfactory as a solution to the problem of unconceived alternatives. Prosecutors and investigators acting epistemically responsibly is neither a necessary nor a sufficient condition for being justified that we have uncovered all plausible alternatives. To see why, consider two fictional cases:

The lazy detective: Detective A is lazy, biased and stuck in her ways. She displays few if any of the epistemic virtues that we’d desire in an investigator. Nonetheless, in the case she is currently working on – the robbery of a jewelry store – she is lucky. She arrives at the crime scene and the store’s employees have already placed the alleged robber under citizen’s arrest. The suspect

²³ Note that on Amaya’s account, these epistemic duties apply to the prosecution and investigators. We may distinguish her proposal from the responsibilist position advocated by Picinali (2015) who interprets the beyond a reasonable doubt standard in terms of epistemically responsible behaviour on the part of the jurors.

²⁴ Admittedly, I am not completely certain that this is the correct interpretation of Amaya’s position. In her writing on the bad lot problem she seems to require *actual* responsible behavior from the relevant agent, as she writes about IBE being warranted only if “one’s set of alternative explanations has been constructed in an epistemically responsible way” (Amaya, 2009, pp. 154–5). However in other writing she adopts a *counterfactual* responsibilist position (Amaya, 2008; 2015). This means that a belief is justified if it *could have been* the outcome of an epistemically responsible process, even if the actual agent holding the belief did not act epistemically responsible. I am inclined to read her position in the former way not only because it most closely fits with how she presents it herself when talking about the bad lot problem, but also because she is not clear on how we ought to assess whether the belief could have been the result from an epistemically responsible process. The counterfactual reading of her work is therefore more vague than the reading I adopt here.

immediately confesses. Finally, the employees show the detective camera footage of him threatening the employees and stealing the store's jewelry. After sloppily taking statements, the detective goes to get a hamburger. She justifies her lack of a proper investigation by telling herself that 'he obviously did it, no need to look further.'

The virtuous detective: Detective B, in contrast, is a hard-working investigator who approaches each case with cleverness and objectivity. However, the current case that she is working on – a murder – proves to be especially complex. It features a host of conflicting witness statements, potentially misleading traces and a towering stack of notes about possible leads. The detective does have a primary suspect who seems to have a means, motive and opportunity. She has also tried hard to come up with a plausible alternative explanation, but failed. However, despite her hard work, she still has the feeling that there may have been something that she missed, as she does not yet see all the connections between the different aspects of the case.

I believe that both detectives are right. The lazy detective justifiably believes that she did not miss any plausible alternatives, despite her lack of a virtuous investigation. After all, the case was straightforward. Any scenario in which the defendant is innocent would have to somehow imply that the testimony of both the employees and the defendant himself, as well as the camera evidence are all misleading. There is no reason to believe a scenario that implies such misleadingness could be plausible. Hence, a responsible investigation is not a necessary condition for justifiably believing that one's set of explanations is exhaustive. In the second case, the investigation *was* responsible. Nonetheless, the detective is not justified in presuming exhaustiveness; she has good reasons to suspect that that she missed some plausible alternatives. So, a virtuous investigation is also not a sufficient condition for the presumption of exhaustiveness to be justified. Responsibilism is therefore inadequate as a solution to the problem of unconceived alternatives.²⁵

²⁵ Note that my examples relate to the detective's belief being justified. Nonetheless, the fact-finder—i.e. the judge or jury—will often base their decision on the set of explanations constructed by investigators during the preceding investigation. Their belief in exhaustiveness will then be similarly (un)justified. For the sake of completeness, let me also mention that responsibilists such as Amaya are often concerned with the *justification of a belief* rather than whether a person is *justified in holding that belief*. They note that there are cases where the two diverge—i.e. where a belief is justified, but the person holding it is not justified in holding it. It is for this reason that Amaya develops the counterfactual position mentioned in the previous footnote. I have no objection against this project. However, in this article I am concerned with when the fact-finder may justifiably believe that her set of explanations is exhaustive. I am not completely certain whether this is also the question that Amaya (2009) tackles when she discusses the bad lot problem. Nonetheless, I believe that whether the fact-finder is justified in presuming exhaustiveness is the appropriate question to ask when talking about the problem of unconceived alternatives. It is, after all, the fact-finder who decides whether the set of scenarios is good enough and I am interested in the rationality of this decision. What the above examples show is that responsibilism is an inadequate answer to this question.

6.2 The Evidentialist View

If not responsibilism, then what? I believe that Amaya rejects the evidentialist option too quickly. On my view, whether we may presume that we have considered all plausible alternatives depends on our evidence. That the underlying investigation displayed certain epistemic virtues is one important (if not the most important) type of evidence for this, but not the only one. For instance, while the lazy detective case featured some evidence in favor of unconceived scenarios existing (the unvirtuous investigation), the evidence as a whole supports the presumption that no plausible alternatives exist. In contrast, in the second case the available evidence does give good reasons to suspect that she may have missed some plausible alternative, even if her investigation was virtuous.

On the evidentialist view we assess whether our evidence indicates that our investigation missed relevant possibilities. Amaya (2009, p. 155) herself already hints at such an evidentialist view when she writes: “provided that one has conducted a thorough search for other potential explanations *and there is no reason that justifies a further search*, then one is justified (in the sense that matters) in accepting as justified the best explanatory hypothesis of the events at trial”.²⁶ This comes close to what I am proposing. Whether a further search is justified or not not only depends on whether the investigation was virtuous, but also on whether we have a reasonable expectation of finding further relevant scenarios, which means considering how likely it is that such alternatives exist. However, I propose something slightly different still. There can also be situations in which we both know that we have likely missed something and that a further search would *not* be likely to produce anything else. Consider the case of the hard-working detective again. Her previous failures to find alternatives might give her an inductive reason to suspect that similar, future endeavors will also not yield anything—that she will not be able to see all the connections between the evidence and thereby come up with all plausible explanations. Nonetheless, she still has good reasons—i.e. good evidence—to suspect that she has missed something.

As I argued before, the question how probable it is that plausible, unconceived alternatives exist is not the same as what the impact of actually discovering such possibilities would be. This latter question will often be impossible to answer because it would require us to know the details of scenarios that we have not even come up with (see Sect. 4). In contrast, the first question is one that we can—at least sometimes—meaningfully reason about. Various philosophers of science write about the kinds of evidence that indicate (an absence of) unconceived alternatives. For example, Stanford (2006) invokes the failure of past scientists to conceive of all relevant theories to argue that contemporary scientists have likely also missed such alternatives. Musgrave (1988) and Lipton (2003) note that, in order to assuage the worry of unconceived alternatives, our best scientific theory should also be ‘satisfactory’ or ‘good enough’. Dawid et al. (2015) suggest that the possibility of unconceived alternatives plausibly depends on the difficulty of the relevant scientific problem, on the

²⁶ My italics.

cleverness of the scientists and on the available resources to investigate the problem. Similarly, Dellsén (2017b) argues that the probability of unconceived alternatives remaining undiscovered depends on the complexity of the domain, how dogmatic and skilled the scientists in that domain are and how long they have been searching for alternatives. Using these ideas as inspiration, I propose the following (putative) criteria for assessing whether the investigation likely uncovered all plausible scenario:

A. *Quality of the investigation*: All other things being equal, the better the search for alternative scenarios the more reason we have to presume that the investigators have uncovered all relevant possibilities. The quality of such a search depends in part on the amount of time and resources spent on it as well as on the imaginative faculties of the investigators. It also depends on the kinds of virtues central to Amaya's account, such as open-mindedness and perseverance.

B. *Quality of the conceived scenarios*: All other things being equal, the better our conceived scenarios, the more reason we have to presume that no plausible alternatives exist. If our conceived scenarios explain everything adequately, then this gives us a good reason to suspect that there are no relevant alternatives. Conversely, if we only have implausible scenarios, then we have good reasons to suspect that there is a better scenario that we have not conceived of. Furthermore, our existing explanations can be of a high quality in the sense that they are specific. It is sometimes easier to think of alternative explanations when we know precisely what we are seeking an alternative for.

C. *Quality of the evidence*: All other things being equal, the better our evidence, the more reason we have to presume that we have uncovered all relevant possibilities. There are at least two ways in which the quality of the evidence matters for whether we have reason to believe that we have may have missed something. First, the evidence that we have in a case steers us in the direction of possible alternative explanations. For instance, footprints not belonging to the main suspect near a crime scene may suggest that there were other possible perpetrators. However, this also means that if the evidence is scarce, such guidance may be limited. For example, if the victim is found in the middle of the forest, with no witnesses, obvious fingerprints or other marks of a possible perpetrator, and if it is not clear whether he was killed or died of natural causes, then it is almost pure guesswork as to what happened. In contrast, if there are multiple witnesses and other evidence that clues us in on what happened, then there is less of a chance that we have overlooked something. However, secondly, a great deal of low-quality evidence can also be misleading as it may create noise, making it harder to determine whether we have overlooked anything.

D. *Difficulty of the case*: All other things being equal, the easier a case, the more reason we have to presume that we have uncovered all relevant possibilities. First, some cases are more complex than others. How complex a case is depends on how difficult it is to 'visualize' what is going on—i.e. to see how all the facts hang together (van Oorschot, 2014). For instance, the case may be accompanied by a thick case file, indicating a great deal of potentially relevant evidence, which may also conflict or be incomplete. It may then be difficult to evaluate how the individual items of evidence should be interpreted or how these items interlink. In such

cases it may be unclear whether all plausible alternatives were considered.²⁷ This is, for instance, why the virtuous detective has good reasons to suspect further unconceived alternatives. Secondly, cases may also be difficult in ways that do not relate to their complexity. For instance, as Amaya (2015, p. 517) points out, in emotionally disturbing cases, investigators may be more likely to be biased, thereby failing to conceive of plausible alternatives.

Let me say something about how these criteria relate to one another. They are *not* necessary and sufficient conditions for being justified in believing our set of scenarios to be exhaustive. Instead, they are items of evidence which *may* jointly justify such a belief. As with any form of evidential reasoning, when a belief in exhaustiveness is justified will depend on the particulars of the case and may involve weighing these factors against one another. Furthermore, the stakes involved can also be relevant. Cases with higher stakes typically require better evidence for our belief to be justified.²⁸

The aforementioned factors also overlap to a degree. For instance, having a great deal of evidence may matter for both the quality of the evidence and the difficulty of the case. Furthermore, how a situation scores on one criterium will often depend on the determination of the other criteria as well. For example, the difficulty of the case determines in part how thorough one's investigation must be. This is why, for example, the lazy detective is justified in their belief that there are no unconceived alternatives—her investigation did not need to be thorough because the case was straightforward. Similarly, as mentioned earlier, whether a scenario offers an adequate explanation of the facts may depend on whether we believe that there are further alternatives to that scenario. Whether this is the case will depend in part on the other criteria.

The final point that I want to make relates back to my rejection of responsibility. Responsibility such as Amaya, but also Baehr (2009) and Cloos (2015) suggest that one's evidence supporting a belief does not always justify one in holding that belief—it may be the case that one's evidence-gathering process was defective. They therefore add a requirement of epistemically responsible inquiry for justification. Above, I rejected responsibility as a solution to the problem of unconceived alternatives. My argument was that whether one acted virtuously in constructing one's set of scenarios is neither a necessary nor a sufficient condition for presuming that this set is exhaustive. Instead I proposed an evidentialist approach. However, suppose that our set of evidence supports the belief that the set of scenarios in a case is exhaustive, but only because one did not spend sufficient time gathering or considering the relevant evidence. If one had been thorough, one would have seen that this belief in exhaustiveness is actually unjustified. In other words, the belief that 'my set of scenarios is exhaustive' differs from the belief that 'my evidence

²⁷ A similar point is made by Roush (2005, pp. 211–2) with respect to scientific theories. She argues that if a domain of inquiry calls for very complicated theories, then it may be harder to conceptualize what alternatives to theories in that domain would look like.

²⁸ For instance, Amaya (2015, 527–8) and Josephson, (2001, 1626) both suggest that the stakes matter for proof beyond a reasonable doubt.

supports the belief that my set of scenarios is exhaustive.’ I reject responsibility with respect to the first belief. However, I want to leave the door open for responsibility with respect to the second (though I do not want to definitively embrace it either). In this sense my proposal here is not at odds with the general responsibility framework at large, only with responsibility with respect to the problem of unconceived alternatives.

7 Conclusion

In criminal cases, reasonable doubts sometimes arise about the defendant’s guilt. For instance, our evidence may be insufficiently strong, the prosecution might fail to present a plausible scenario or the defense can present a plausible alternative scenario. Yet another, mostly overlooked, source of reasonable doubt is the possible existence of plausible, *unconceived* alternatives. There are specific situations in which we have good reasons to presume that we might have failed to imagine some relevant alternative scenario. In such cases, we cannot consider the guilt of the defendant proven. In this article, I discussed this claim through the lens of two influential conceptions of rational criminal proof: the Bayesian and the explanation-based.

What this normative, epistemic claim means for the practice of criminal proof will depend on the particulars of the legal system. One option is that the problem of unconceived alternatives imposes a burden on the prosecution to prove that all relevant possibilities were considered. On the other hand, we could say that it is up to the defense to raise possibility of unconceived alternatives in specific situations, to argue that there is a reasonable doubt about the defendant’s guilt. Perhaps the burden to ensure that our set of scenarios is good enough should be placed at the feet of investigators. Alternatively, Nance (2008) suggests that the decision whether the available set of evidence is sufficiently complete for the case to go to trial lies with the judge. Whether we suspect unconceived alternatives could be part of this decision. Finally, the answer might be ‘all of the above’.

Acknowledgements I would like to thank Anne Ruth Mackor, Henry Prakken, Pepa Mellema, Job de Grefte, Astrid Voorwinden, Jan-Willem Romeijn, Kritika Maheshwari, Finnur Dellsén and the four anonymous reviewers for their helpful comments on previous versions of this article.

Funding This work is part of the research programme with Project Number 160.280.142, which is (partly) financed by the Netherlands Organisation for Scientific Research (NWO).

Declaration

Competing interests The author declares none.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission

directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Allen, R. J. (2010). No plausible alternative to a plausible story of guilt as the rule of decision in criminal cases. In L. Laudan & J. Cruz (Eds.), *Proof and standards of proof in the law* (pp. 10–27). Northwestern University School of Law.
- Allen, R. J., & Pardo, M. S. (2007). Probability, explanation and inference: A reply. *The International Journal of Evidence & Proof*, 11(4), 307–317.
- Allen, R. J., & Pardo, M. S. (2017). Explanations and the preponderance standard: Still kicking rocks with Dr. Johnson. *Seton Hall Law Review*, 48, 1579.
- Allen, R. J., & Pardo, M. S. (2019). Relative plausibility and its critics. *The International Journal of Evidence & Proof*, 23(1–2), 5–59.
- Amaya, A. (2015). *The tapestry of reason: An inquiry into the nature of coherence and its role in legal argument*. Bloomsbury Publishing.
- Amaya, A. (2008). Justification, coherence, and epistemic responsibility in legal fact-finding. *Episteme*, 5(3), 306–319.
- Amaya, A. (2009). Inference to the best legal explanation. In H. Prakken, H. Kaptein, & B. Verheij (Eds.), *Legal evidence and proof: Statistics, stories, logic* (pp. 135–160). Ashgate Publishing.
- Baehr, J. (2009). 'Evidentialism, Vice, and Virtue'. *Philosophy and Phenomenological Research*, 78(3), 545–567.
- Bex, F. J. (2011). *Arguments, stories and criminal evidence: A formal hybrid theory* (Vol. 92). Springer Science & Business Media.
- Bex, F., & Walton, D. (2012). Burdens and standards of proof for inference to the best explanation: Three case studies. *Law, Probability and Risk*, 11(2–3), 113–133.
- Biedermann, A., & Vuille, J. (2018). The decisional nature of probability and plausibility assessments in juridical evidence and proof. *International Commentary on Evidence*, 16(1), 1–30.
- Blackstone, W. (1962). *Commentaries on the Laws of England* (Vol. 2). Beacon.
- Bradley, R. (2017). *Decision theory with a human face*. Cambridge University Press.
- Cheng, E. K. (2013). Reconceptualizing the burden of proof. *Yale LJ*, 122, 1254–1279.
- Cloos, C. M. (2015). Responsibilist evidentialism. *Philosophical Studies*, 172(11), 2999–3016.
- Dahlman, C. (2019). De-biasing legal fact-finders with Bayesian thinking. *Topics in Cognitive Science*.
- Dawid, R., Hartmann, S., & Sprenger, J. (2015). The no alternatives argument. *The British Journal for the Philosophy of Science*, 66(1), 213–234.
- Dellsén, F. (2017a). Reactionary responses to the bad lot objection. *Studies in History and Philosophy of Science Part A*, 61, 32–40.
- Dellsén, F. (2017b). Realism and the absence of rivals. *Synthese*, 194(7), 2427–2446.
- Dellsén, F. (2018). The heuristic conception of inference to the best explanation. *Philosophical Studies*, 175(7), 1745–1766.
- Douven, I. (2017). What is inference to the best explanation? And Why Should We Care? In T. Poston & K. McCain (Eds.), *Best explanations: new essays on inference to the best explanation* (pp. 431–470). Oxford University Press.
- Dworkin, R. (1985). *A matter of principle*. OUP Oxford.
- Earman, J. (1992). *Bayes or Bust? A critical examination of bayesian confirmation theory*. Cambridge: MIT Press.
- Eppe, D. (2014). The consequences of error in criminal justice. *Harvard Law Review*, 128, 1065.
- Fenton, N., Lagnado, D., Dahlman, C., & Neil, M. (2019). The opportunity prior: A proof-based prior for criminal cases. *Law, Probability and Risk*, 18(4), 237–253.
- Fenton, N., Neil, M., & Berger, D. (2016). Bayes and the Law. *Annual Review of Statistics and Its Application*, 3, 51–77.
- Fitelson, B., & Thomason, N. (2008). Bayesians sometimes cannot ignore even very implausible theories (even ones that have not yet been thought of). *The Australasian Journal of Logic*, 6.
- Gardiner, G. (2019). The Reasonable and the Relevant: Legal Standards of Proof. *Philosophy & Public Affairs*, 47(3), 288–318.

- Gillies, D. (2001). Bayesianism and the fixity of the theoretical framework. In *Foundations of Bayesianism* (pp. 363–379). Springer.
- Goldman, A. I. (2001). Quasi-objective Bayesianism and legal evidence. *Jurimetrics*, 42, 237.
- Hamer, D. A. (2004). Probabilistic standards of proof, their complements and the errors that are expected to flow from them. *University of New England Law Journal*, 1, 71.
- Ho, H. L. (2015). The legal concept of evidence. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2015 Ed.). <https://plato.stanford.edu/entries/evidence-legal/>
- Jellema, H. (2020). The reasonable doubt standard as inference to the best explanation. *Synthese*.
- Josephson, J. R. (2001). On the proof dynamics of inference to the best explanation. *Cardozo Law Review*, 22, 1621–1643.
- Kuipers, T. A. (2004). Inference to the best theory, rather than inference to the best explanation—Kinds of abduction and induction. In *Induction and deduction in the sciences* (pp. 25–51). Springer.
- Laudan, L. (2006). *Truth, error, and criminal law: An essay in legal epistemology*. Cambridge University Press.
- Laudan, L. (2007). Strange bedfellows: Inference to the best explanation and the criminal standard of proof. *The International Journal of Evidence & Proof*, 11(4), 292–306.
- Lipton, P. (2003). *Inference to the best explanation*. Routledge.
- Mackor, A. R., Jellema, H., & van Koppen, P. L. (2021). Explanation-based approaches to reasoning about evidence and proof in criminal trials. In B. Brozek & J. Hage (Eds.), *Law and mind: A survey of law and the cognitive sciences* (pp. 431–470). Cambridge University Press.
- Meester, R., & Sjerps, M. (2004). Why the effect of prior odds should accompany the likelihood ratio when reporting DNA evidence. *Law, Probability and Risk*, 3(1), 51–62.
- Musgrave, A. (1988). The ultimate argument for scientific realism. In *Relativism and realism in science* (pp. 229–252). Springer.
- Nance, D. A. (2008). The weights of evidence. *Episteme: A Journal of Social Epistemology*, 5(3), 267–281.
- Pardo, M. S., & Allen, R. J. (2008). Juridical proof and the best explanation. *Law and Philosophy*, 27(3), 223–268.
- Pennington, N., & Hastie, R. (1993). The story model for juror decision making. In R. Hastie (Ed.), *Inside the juror: The psychology of juror decision making* (pp. 192–221). New York: Cambridge University Press.
- Picinali, F. (2015). The threshold lies in the method: Instructing jurors about reasoning beyond reasonable doubt. *The International Journal of Evidence & Proof*, 19(3), 139–153.
- Ribeiro, G. (2018). Can there be a burden of the best explanation? *The International Journal of Evidence & Proof*, 22(2), 91–123.
- Roberts, P. (2020). Presumptuous of pluralistic presumptions of innocence? Methodological diagnosis towards conceptual reinvigoration. *Synthese*.
- Roush, S. (2005). *Tracking truth: Knowledge, evidence, and science*. Clarendon Press.
- Rowbottom, D. P. (2019). Extending the argument from unconceived alternatives: Observations, models, predictions, explanations, methods, instruments, experiments, and values. *Synthese*, 196(10), 3947–3959.
- Salmon, W. (1990). Rationality and objectivity in science or Tom Kuhn meets Tom Bayes. *Scientific Theories*, 14, 175–204.
- Sklar, L. (1981). Do unborn hypotheses have rights? *Pacific Philosophical Quarterly*, 62, 17–29.
- Stanford, P. K. (2006). *Exceeding our grasp: Science, history, and the problem of unconceived alternatives* (Vol. 1). Oxford University Press.
- Steele, K., & Stefánsson, H. O. (2019). Belief revision for growing awareness. *Mind*.
- Stein, A. (2005). *Foundations of evidence law*. New York University Press.
- Strong, J. W., & Broun, K. S. (1992). *McCormick on evidence*. West Publishing Company.
- Talbott, W. (2008). Bayesian epistemology. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Fall 2007 ed.). <https://plato.stanford.edu/entries/epistemology-bayesian/>
- Urbaniak, R. (2018). Narration in judiciary fact-finding: A probabilistic explication. *Artificial Intelligence and Law*, pp. 1–32.
- Urbaniak, R., & Di Bello, M. (2021). Legal Probabilism. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Fall 2021 ed.). <https://plato.stanford.edu/entries/legal-probabilism/>
- van Koppen, P. J. (2011). *Overtuigend bewijs: Indammen van rechterlijke dwalingen [Convincing evidence: Reducing the number of miscarriages of justice]*. Nieuw Amsterdam.
- Van Fraassen, B. C. (1989). *Laws and symmetry*. Oxford University Press.

- Van Oorschot, I. (2014). Seeing the case clearly: File-work, material mediation, and visualizing practices in a Dutch criminal court. *Symbolic Interaction*, 37(4), 439–457.
- Vlek, C., Prakken, H., Renooij, S., & Verheij, B. (2013). Modeling crime scenarios in a Bayesian network. In: *Proceedings of the fourteenth international conference on artificial intelligence and law*, pp. 150–159.
- Wenmackers, S., & Romeijn, J.-W. (2016). New theory about old evidence. *Synthese*, 193(4), 1225–1250.

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