Reconciling Bayesian Epistemology and Narration-based Approaches to Judiciary Fact-finding

Rafal Urbaniak

Centre for Logic and Philosophy of Science, Ghent University, Belgium[*]

Institute of Philosophy, Sociology and Journalism, University of Gdansk, Poland[†]

rfl.urbaniak@gmail.com

Legal probabilism (LP) claims the degrees of conviction in juridical fact-finding are to be modeled exactly the way degrees of beliefs are modeled in standard bayesian epistemology. Classical legal probabilism (CLP) adds that the conviction is justified if the credence in guilt given the evidence is above an appropriate guilt probability threshold. The views are challenged on various counts, especially by the proponents of the so-called narrative approach, on which the fact-finders' decision is the result of a dynamic interplay between competing narratives of what happened. I develop a way a bayesian epistemologist can make sense of the narrative approach. I do so by formulating a probabilistic framework for evaluating competing narrations in terms of formal explications of the informal evaluation criteria used in the narrative approach.[1]

1 Legal probabilism and its troubles

According to *legal probabilism* (LP), degrees of conviction in juridical fact-finding are to be modeled exactly the way degrees of beliefs are modeled in standard bayesian epistemology: by means of probabilistic distributions satisfying the standard axioms of probability theory. *Classical legal probabilism* (CLP), which originated with Bernoulli [3] adds on top of that the view according to which the criminal standard of proof beyond reasonable doubt should be equated with a certain high threshold probability of guilt (albeit, some variants of the view admit that thresholds for different cases might be different).[2]

LP (and generally, the use of probabilistic methods in judiciary contexts) is criticized from various angles (see for example [18], [19], [4], [20], [14], [5], [6], [21], [17], [2], [10], [9]). The critics of LP argue that the view is blind to various phenomena that an adequate philosophical account of legal fact-finding should explain. Some of them pertain to procedural issues [17]: proceedings are back-and-forth between opposing parties, cross-examination is crucial, and yet CLP seems to take no notice of this dynamics. Some have to do with reasoning methods which are not only evidence-to-hypothesis, but also hypotheses-to-evidence [21, 2] and involve inference to the best explanation [6]. A better account, arguably, is one in which the proceedings are seen as an interplay of evidence and various explanations (often called *narratives*) presented by opposing parties [10].

Accordingly, the no plausible alternative story (NPAS) theory [1] is that the courtroom is a confrontation of competing narrations offered by the defendant and by the prosecutor and the narrative to be selected should be the most plausible one. The view is conceptually plausible [7] and finds support in

---

[1]Due to the nature of this volume and space restrictions, this paper focuses mostly on presenting the technical apparatus. Deeper philosophical discussion, application to examples, and further informal clarifications are postponed to a different, lengthy paper.

[2]Nowadays, legal scholars fond of probabilism usually inscribe to LP but not to CLP.

psychological evidence [16, 15]. The approach is also better at capturing the already listed phenomena that CLP is claimed to be blind to.

From the perspective of a formal epistemologist, the key disadvantage of NPAS is that it abandons the rich toolbox of probabilistic methods and takes the key notion of plausibility to be a primitive notion which should be intuitively understood.[3] The goal of this paper is to develop a bayesian approach to NPAS, showing that one can embrace NPAS without giving up on probabilistic methods.

2   Classical legal probabilism and the gatecrasher paradox

CLP is also susceptible to criticism directed at its specific claim: that there is a guilt probability threshold. One of the most well-known conceptual arguments against the idea is the following.

The *paradox of the gatecrasher* [4][4] has been developed to indicate that mere high probability of guilt is not enough for a conviction. A variant of the paradox goes as follows:

> Suppose our guilt threshold is high, say at 0.99. Consider the situation in which 1000 fans enter a football stadium, and 999 of them avoid paying for their tickets. A random spectator is tried for not paying. The probability that the spectator under trial did not pay exceeds 0.99. Yet, intuitively, a spectator cannot be considered guilty on the sole basis of the number of people who did and did not pay.

CLP fails to handle the gatecrasher paradox: whatever $\neq 1$ threshold you propose, I can give you a variant of the gatecrasher with sufficiently many people for the guilt statement of an arbitrary spectator to be above the guilt threshold.

NPAS, on the other hand, seems to avoid the paradox: conviction of a random attendee is unjustified, because there is no real plausible story of guilt of that particular person. After all, merely quoting the statistics, on this approach, hardly counts as giving a narrative. The problem is, such considerations are left at an informal level, the requirements on what should count as a narration aren't too clear, and so the considerations remain somewhat indecisive. After all, a stubborn philosopher might still insist: *look, all that is needed is a theory of what happened that is more plausible than the alternative — my theory is, the suspect is guilty, and, given the evidence, it is more plausible than the alternative, which says that he isn't.*

3   New legal probabilism

New legal probabilism (NLP) is an attempt to improve on the underspecificity of NPAS [7]. While still being informal, the approach is more specific about the conditions that a successful accusing narration is to satisfy for the conviction beyond reasonable doubt to be justified. Di Bello identifies four key requirements.

---

[3]"I have been asked to elaborate on the meaning of plausibility in this theoretical structure. The difficulty in doing so is that the relative plausibility theory is a positive rather than a normative theory. For reasons I will elaborate below, plausibility can serve as a primitive theoretical term the meaning of which is determined by its context in the explanation of trials." [1, p. 10]

[4]It is analogous the *prisoners in a yard scenario* [14], where a group of prisoners commits a group killing, and it's impossible to identify the single innocent prisoner. Mathematically, the examples are pretty much the same.

(Evidential support)	The defendant's guilt probability on the evidence should be sufficiently supported by the evidence, and a successful accusing narration should explain the relevant evidence.
(Evidential completeness)	The evidence available at trial should be complete as far as a reasonable fact-finders' expectations are concerned.
(Resiliency)	The prosecutor's narrative, based on the available evidence, should not be susceptible to revision given reasonably possible future arguments and evidence.
(Narrativity)	The narrative offered by the prosecutor should answer all the natural or reasonable questions one may have about what happened, given the content of the prosecutor's narration and the available evidence.

How would NLP handle the gatecrasher? Di Bello's discusses an analogous *prisoners* scenario, in which a footage indicates that all but one prisoners participated in a group killing, but the footage doesn't allow for identification of the innocent prisoner [7, pp. 219-222]. He argues that the conviction is unjustified on two counts. First, the narrative is grossly incomplete: "What is the initiating event? What psychological responses did it trigger? Who participated in the killing? What were the different participants doing?..." Second, the convicting scenario, Di Bello argues, isn't resilient, for it is quite plausible that more evidence might become available.

Unfortunately, it is not quite clear whether the resources that NLP helps herself to when formulating the requirements, indeed fall into the realm of bayesian methodology:

> The probabilists can enrich their framework by adding probability-based accounts of evidential completeness, resiliency, and narrativity. To my knowledge, no legal probabilist has undertaken the task in any systematic way. [7, p. 75]

In what follows we undertake this task.

4 Preliminary technicalities

4.1 The language and its interpretation

The *object language* is a standard propositional language $\mathscr{L}$ (I assume $\wedge$ and $\neg$ are the primitive connectives, nothing serious hinges on the choice) extended to a language $\mathscr{L}^+$ with primitive unary operators $E, N_1^A, \ldots, N_k^A, N_1^D, \ldots, N_k^D$, and the guilt statement constant G.

The content of the guilt statement is given in terms of a list of conditions in the background language that need to be established for a conviction to be justified. This is modeled by conditioning on the definition of guilt G which has the form of $G \equiv g_1 \wedge \cdots \wedge g_l$ for appropriate $g_1, \ldots, g_l \in \mathscr{L}$.

The intended interpretation of Ep is *p is part of evidence*. The idea is that after all the evidence and all the arguments have been presented in court, the background knowledge is to be enriched by the pieces of evidence presented, thought of as sentences of $\mathscr{L}$: $E = \{e_1, \ldots, e_j\} \subset \mathscr{L}$. However, we are not only to extend our beliefs by $e_1, \ldots, e_j$, but also by the corresponding claims about these sentences being part of evidence: $Ee_1, \ldots, Ee_j$.

$\mathbb{N}_i^A p$ means *p is part of an accusing narration* $\mathbb{N}_i^A$ and $\mathbb{N}_i^D p$ means *p is part of a defending narration* $\mathbb{N}_i^D$. Each narration $\mathbb{N}_i$ (in contexts in which it is irrelevant whether a narration is an accusing one or not, I will suppress the superscripts) is taken to be a finite set of sentences $n_{i1}, n_{i2}, \ldots, n_{ik_i}$ of $\mathscr{L}^+$.

E stands ambiguously for the set of all sentences constituting evidence and for the conjunction thereof. Which reading is meant will always be clear from the context (this convention applies to all finite sets of sentences considered in this paper). $\mathbb{E}^d$ stands for $\mathbb{E}^d = \{E\varphi | \varphi \in \mathbb{E}\}$ and $\mathbb{E}^-$ is $\{\neg E\varphi | \varphi \notin \mathbb{E}\}$. The distinction is needed, because there is a difference between knowing that certain sentences are pieces of evidence, and knowing that no other sentence is. For any narration $\mathbb{N}_i$, symbols $\mathbb{N}_i$, $\mathbb{N}_i^d$, and $\mathbb{N}_i^-$ are to be understood analogously to E, $\mathbb{E}^d$ and $\mathbb{E}^-$. $\mathbb{N}^d$ is the (positive) *description* of all the narrations, $\bigcup_i \mathbb{N}_i^d$, and $\mathbb{N}^- = \bigcup_i \mathbb{N}_i^-$ adds that this description is complete.

4.2   Partial probability and four thresholds

The current framework diverges from the standard bayesianism in using partial probability functions rather than full probabilistic discribtions to model credences. This is motivated by noticing that the fact-finders, on one hand, are supposed to rely on their background knowledge when assessing the plausibility of a given narration, but on the other hand, they clearly cannot rely on all biases and assumptions that they have.[5]

A *partial conditional credence function* P (partially) maps $\mathscr{L}^+ \times \mathscr{P}(\mathscr{L}^+)$ to $[0,1]$.[6] (I often write $P(h|\mathbb{E})$ instead of $P(\langle h, \mathbb{E} \rangle)$). Let $\downarrow$ and $\uparrow$ stand for *being defined* and *being undefined* respectively. A partial probability distribution has to have an extension to a total conditional probability distribution over $\mathscr{L}^+$ satisfying the standard axioms of conditional probability, Moreover, it additionally has to satisfy the following conditions for any $\Gamma \subseteq \mathscr{L}^+$, and any $\varphi, \psi \in \mathscr{L}^+$:

$$P(\top|\Gamma) = 1 \qquad\qquad P(\bot|\Gamma) = 0 \qquad\qquad \text{(Part-1)}$$

$$\varphi \in \Gamma \Rightarrow P(\varphi|\Gamma)\downarrow \qquad\qquad\qquad\qquad\qquad\qquad \text{(Part-2)}$$

$$P(\varphi|\Gamma)\downarrow \Leftrightarrow P(\neg\varphi|\Gamma)\downarrow \qquad P(\varphi \wedge \psi|\Gamma)\downarrow \Leftrightarrow P(\psi \wedge \varphi|\Gamma)\downarrow \qquad \text{(Part-3)}$$

$$P(\varphi \wedge \psi|\Gamma) > 0 \Rightarrow P(\varphi|\Gamma)\downarrow, P(\psi,|\Gamma)\downarrow \qquad P(\varphi|\Gamma) = 0 \Rightarrow P(\varphi \wedge \psi|\Gamma)\downarrow \qquad \text{(Part-4)}$$

$$P(\varphi|\Gamma)\uparrow \Rightarrow P(\varphi \wedge \psi|\Gamma)\uparrow \qquad\qquad \text{unless } P(\psi|\Gamma) = 0 \qquad \text{(Part-5)}$$

$$\text{If } P(\varphi|\Gamma) > 0, P(\varphi \wedge \psi|\Gamma) = 0, \qquad\qquad \text{then } P(\psi|\Gamma) = 0 \qquad \text{(Part-6)}$$

(Part-1) requires that logical truths have probability 1 and logical contradictions always have probability 0. (Part-2) requires that the probability of a claim given a set of premises that includes it is always defined. (Part-3) states that the conditional probability of a claim is defined just in case the conditional probability of its negation is, and that the order of conjuncts has no impact on whether the conditional probability of a conjunction is defined. According to (Part-4), the conditional non-zero probability of a conjunction is defined only if the conditional probability of both conjuncts is. Moreover, if the conditional probability of a conjunct is 0, the conditional probability of the conjunction is defined (and by the fact that the credence has an extension to a total conditional probability satisfying the standard axioms, we also know that it will be 0 as well). (Part-5) says that, unless this unusual circumstance occurs, the

---

[5]Suspending our conviction about *p* cannot be easily modeled in the standard bayesian framework, since even the most sensible candidate, $1/2$, doesn't do the job. Just to give a simple example, there is a difference between knowing that a given coin is fair and assigning probability of $1/2$ to heads, and not knowing how fair a coin is at all and assigning probability of $1/2$ to heads for this reason.

[6]Partial credence functions for conditional probabilities have been introduced in [13] and [12]; my definition differs from that account in a few inessential aspects.

conditional probability of a conjunction is undefined if the conditional probability of at least one conjunct is. Finally, (Part-6) demands that if a conjunct has a conditional probability > 0, then the conjunction has conditional probability 0 only if the other conjunct does.

Since the fact-finders are supposed not to be biased and aren't informed about the trial yet, we additionally assume that the priors of guilt, of what the evidence and what the narrations are, are undefined: $P(G) \uparrow, P(g_1 \wedge \cdots \wedge g_l) \uparrow, P(E\varphi) \uparrow, P(N_i\varphi) \uparrow$ for any $\varphi \in \mathscr{L}^+$, and any $1 < i < k$.

Four types of stances that a fact-finder might take towards a claim will be considered. First, a fact-finder might consider a claim completely uncontroversial, and accept it without any further argument. Such stance will be modeled by the credence in a given claim reaching the *uncontroversial acceptability threshold*, a. On the opposite side of spectrum we have the *negligibility threshold*, $n =_{df} 1 - a$. Notice that a and n can't be respectively 1 and 0, for this would require complete unrevisable certainty.

One more type of stance needs to be incorporated into the framework — that of *strong plausibility*. Usually there are claims that are strongly supported while not being as close to certainty as the uncontroversially acceptable ones. This kind of credence that we would normally find sufficient for acting upon in our uncertain world will be denoted by s. The opposite will be called *rejectability*, $r =_{df} 1 - s$. Clearly we should require $a > s > r > n$.

4.3   Information and updates

After all the evidence and all the arguments have been presented in court, the background knowledge obtained now consists of the pieces of evidence presented, E, information about what is not part of evidence, E^-, the content of the guilt statement G, and the description of the content of a certain finite assembly of finite theories meant to defend or accuse the defendant, $N_1^D, \ldots, N_k^D, N_1^A, N_m^A$ (and what is not part of which narration). When making various assessments in the fact-finding process, at various stages one needs to conditionalize on various parts of the available information, depending on what is being assessed. The variants used are listed in the table below.

name	notation	meaning
full	$P^f(\varphi\|\Gamma)$	$P(\varphi\|E, E^d, E^-, N^d, N^-, G, \Gamma)$
n-full	$P^{nf}(\varphi\|\Gamma)$	$P(\varphi\|E, E^d, N^d, N^-, G, \Gamma)$
informed	$P^i(\varphi\|\Gamma)$	$P(\varphi\|E, E^d, N^d, G, \Gamma)$
evidential	$P^e(\varphi\|\Gamma)$	$P(\varphi\|E, E^d, E^-, G, \Gamma)$
argued	$P^a(\varphi\|\Gamma)$	$P(\varphi\|N^d, G, \Gamma)$
play-along	$P^{N_j}(\varphi\|\Gamma)$	$P(\varphi\|N_j, N^d, N^-, G, \Gamma)$
n-extended play-along	$P^{nN_j}(\varphi\|\Gamma)$	$P(\varphi\|N_j, E, E^d, N^d, N^-, G, \Gamma)$
e-extended play-along	$P^{eN_j}(\varphi\|\Gamma)$	$P(\varphi\|N_j, E, E^d, E^-, N^d, G, \Gamma)$
f-extended play-along	$P^{fN_j}(\varphi\|\Gamma)$	$P(\varphi\|N_j, E, E^d, E^-, N^d, N^-, G, \Gamma)$

5   Defining conditions on a set of narrations

Think of N as the set of attacking and defending narrations seriously considered in the fact-finding process. In this section I list some basic conditions on a set of sets of sentences to count as a set of narration.

In the next section, I explicate the requirements used in the evaluation of narrations.

$$P^f(\neg(N_i \wedge N_j)) \geq \mathtt{a}, \text{ for } i \neq j \qquad\qquad \text{(Exclusion)}$$

$$P^{fN_i^A}(G) \geq \mathtt{a} \wedge P^{fN_k^D}(\neg G) \geq \mathtt{a} \qquad\qquad \text{(Decision)}$$

$$P^e(N_k) \geq \mathtt{n} \qquad\qquad \text{(Initial plausibility)}$$

$$P^f(N_1 \vee \cdots \vee N_k) \geq s \qquad\qquad \text{(Exhaustion)}$$

(Exclusion) requires that narrations under consideration should pairwise exclude each other given what we know about the case. According to (Decision), a defense narration should clearly state that, given all that is known, the accused is not guilty, and an accusing narration should clearly state that, given all that is known, they are. (Initial plausibility) says that we shouldn't consider narrations that are uncontroversially excluded by sensible background knowledge or by evidence. (Exhaustion) requires that it should be strongly plausible that at least one of the narrations hold, given all that we know about the case.

6 Evaluation criteria

Explaining evidence. Now we're ready to look at the explication of the criteria involved in the evaluation of competing narratives. Let's start with the requirement that they should explain the evidence. After all the narrations have been presented and deployed, an accusing narration N_i^A should "make sense" of evidence in the following sense. For any item of evidence presented, e, if, according to N_i^A, it is not excluded as evidence, it should be strongly plausible given N_i^A.

$$\text{For any } e \in \mathtt{E}, \left[\neg P^{N_i^A}(\neg Ee) \geq \mathtt{s} \Rightarrow P^{N_i^A}(e) \geq \mathtt{s}\right] \qquad \text{(Explaining evidence A)}$$

The sense in which a defending narration is supposed to explain evidence is somewhat different. After all, if the defense story is rather minimal and mostly constitutes in rebutting the accusations, it isn't reasonable to expect the defense to explain all pieces of evidence, as long as they aren't really used to support the opposing accusing narration. Rather, the defense should argue that the possibility of the evidence being as it is while the defense's narration is true hasn't been rejected. Thus, we put the condition on a defending narration N_k^D as follows:

$$\text{For any } e \in \mathtt{E}, \text{ if there is } N_i^A \qquad\qquad \text{(Explaining evidence D)}$$
$$\text{such that } P(N_i^A|e) > P(N_i^A),$$
$$\text{then } P^{N_k^D}(e) \geq \mathtt{r}.$$

Missing evidence. The intuition here is that sometimes, given the narration and whatever evidence we already have, certain evidence should be available, but it isn't. For instance, in a drunk driving case the fact-finders would naturally expect a breathalyzer result, and in a murder case evidence as to how the victim was killed is needed.

$$\mathbf{ME}(N_i) \Leftrightarrow \text{for some } \varphi_1, \ldots, \varphi_u \notin \mathtt{E}: \qquad\qquad \text{(Missing evidence)}$$
$$\left[P^{nN_i}(E(\varphi_1) \vee \cdots \vee E(\varphi_u)) \geq \mathtt{s}\right]$$

The disjunction above is there to ensure generality: it might be the case that some evidence from among a group of possible pieces of evidence would be needed, without any particular piece of evidence being expected.

Gaps. Sometimes a narration should be more specific, given what it says and what we already know. For instance, an accusing narration might be required to specify how the victim was attacked, or a defending narration should specify where the defendant was at the time of the crime. Accordingly, we say that $\mathtt{N}_i$ is **gappy** ($\mathbf{G}(\mathtt{N}_i)$) just in case there are claims that the narration should choose from and yet it doesn't:

$$\mathbf{G}(\mathtt{N}_i) \Leftrightarrow \text{for some } \varphi_1, \ldots, \varphi_u \notin \mathtt{N}_i \qquad\qquad \text{(Gap)}$$
$$\mathtt{P}^{f\mathtt{N}_i}(\varphi_1 \vee \cdots \vee \varphi_u) \geq \mathtt{s} \wedge$$
$$\mathtt{P}^{e\mathtt{N}_i}(N_i(\varphi_1) \vee \cdots \vee N_i(\varphi_u)) \geq \mathtt{s}$$

Dominating accusing narration. An accusing narration $\mathtt{N}_i^A$ *dominates* the set of all accusing narrations $\mathbb{N}^A$ just in case it doesn't miss any evidence, it doesn't contain any gap, in light of all available information and evidence it is at least as likely any other accusing narration, and it is strongly plausible, given all available information:

$$\mathbf{D}(\mathtt{N}_i^A) \Leftrightarrow \neg\mathbf{ME}(\mathtt{N}_i^A) \wedge \neg\mathbf{G}(\mathtt{N}_i^A) \wedge \qquad\qquad \text{(Domination)}$$
$$\mathtt{P}^f(\mathtt{N}_i^A) \geq \mathtt{P}^f(\mathtt{N}_j^A) \text{ for all } j \neq i \wedge$$
$$\mathtt{P}^f(\mathtt{N}_i^A) \geq \mathtt{s}$$

Resiliency. A dominating narration $\mathtt{N}_i^A$ is *resilient* ($\mathbf{R}(\mathtt{N}_i^A)$) just in case there is no non-negligible potential evidence that might undermine it, at least in light of all we know (minus the negative description of the evidence, to avoid triviality) — that is, no φ with $\mathtt{P}^{nf}(E\varphi) \geq \mathtt{n}$ – such that if E was modified to $\mathtt{E} \cup \{\varphi\}$, $\mathtt{N}_i^A$ would no longer dominate.

Conviction beyond reasonable doubt. A defense narration $\mathtt{N}_k^D$ *raises reasonable doubt* ($\mathbf{RD}(\mathtt{N}_k^D)$) if it has no gaps, and hasn't been rejected given all that we know:

$$\mathbf{RD}(\mathtt{N}_k^D) \Leftrightarrow \neg\mathbf{G}(\mathtt{N}_k^D) \wedge \mathtt{P}^f(\mathtt{N}_k^D) \geq \mathtt{r} \qquad\qquad \text{(Reasonable doubt)}$$

Accordingly, we say that a conviction is *beyond reasonable doubt* if it is justified by a resilient dominating narration and no defense narration raises reasonable doubt.

7   Looking at the gatecrasher paradox

Once we've formulated the framework, it will be instructive to use it to look at the gatecrasher paradox, and to observe how the requirements and distinctions introduced help us obtain better insight.

Let's creatively call the suspects 1, 2, 3, ..., 1000. Consider the situation in which the accusing narration is simply *1 gatecrashed*, $\mathtt{g}_1$, and the defending narration is simply *1 didn't gatecrash*, $\neg\mathtt{g}_1$. Suppose we have $\mathtt{P}(\mathtt{g}_1|\mathtt{e}) = 999/1000$ and $\mathtt{P}(\neg\mathtt{g}_1|\mathtt{e}) = 1/1000$.

Variant 1 $\mathbb{N}^A = \mathtt{g}_1, \mathbb{N}^D = \neg\mathtt{g}_1.$

One might try to counter this formulation with the following strategy:

Strategy 1: Extreme thresholds   Take $\mathtt{P}^f(\cdot)$ to be $\mathtt{P}(\cdot|\mathtt{e})$, and claim $\mathtt{r} < 1/1000$ and $\mathtt{s} > 999/1000$.

The strategy isn't too successful, though — it's *ad hoc* and it isn't immune to other versions of the paradox, where the number of people is tweaked to get the guilt probability threshold about whatever threshold was picked in Strategy 1. Perhaps, a better approach would be this:

Strategy 2: Full credence vs. statistics	Take $P^f(\cdot)$ to **not** be $P(\cdot|e)$, and claim $P^f(g_1)$ is insufficiently high.

Perhaps there is something to saying that the posterior credences don't have to match the statistical probabilities (as Kaye [11] suggests). But merely saying so doesn't provide a principled explanation: the fact that we're not inclined to accept a claim that has high statistical probability is the *explanandum*, not the *explanans*.

Strategy 3: realistic gaps and non-resiliency	There are multiple natural questions that the accusing narration fails to answer; it also fails to ensure that no future evidence is likely which might overturn the decision.

This is, pretty much, the strategy pursued in [7]. While it is a perfectly valid strategy if the question is how a case like gatecrasher would be handled in reality, this is not the end of the story.

The gatecrasher, or at least a variant thereof, can be viewed rather as an abstract thought experiment formulated to make a conceptual point. And when a philosopher formulates an abstract thought experiment, they're free to set it up any way they want, without too much attention paid to the level of realism, as long as appropriate conceptual restraints are satisfied.

In particular, one might simply make it part of the description of the situation that no further evidence can be obtained despite everyone's best attempts, and no realistic considerations about further details can play a role. No witnesses can come and testify as to the character of 1, because he's been living under a rock and knows nobody. No question about the exact location is natural, because the stadium is somewhat unusual in having no gates, and in fact the 999 people who didn't pay for tickets jumped the fence, in an evenly distributed manner. No surveillance was possible because an accident in the local atomic plant fried all electronics in the area, etc. Call such a variant of the paradox *abstract gatecrasher*.

The question now is: putting the issues with realism and supposing the abstract gatecrasher is immune to these, can we think of more purely epistemic or deontic reasons why the accusing narration in the abstract gatecrasher is not sufficient for a conviction beyond reasonable doubt?

One issue that may come to mind even from a more abstract perspective, is that it doesn't seem that the accusing narration $N_1 = g_1$ explains the evidence.

Strategy 5: Unexplained evidence	The narration relies on e without entailing that it shouldn't be evidence, and yet $P^{N_1}(e) \not\geq s$.

The mere fact that 1 decided to gatecrash, while suggesting that he might've done it with a bunch of friends, certainly doesn't make the claim that so did other 998 people strongly plausible.

This, however, can be fairly easily fixed by modifying the accusing narration to *1 gatecrashed, and so did 998 other people*. This, formally comes to using

$$\boxed{\text{Variant 2} \quad N_1' = \{g_1, e\}}$$

Indeed, the mere fact that 1 decided to gatecrash, while suggesting that he might've done it with a bunch of friends, certainly doesn't make the claim that so did other 998 people strongly plausible. Quite clearly, $P^{N_1}(e|N_1') \geq s$, and so Strategy 5 doesn't beat the paradox with this narration as the accusing narration. Can we do any better?

8   Epistemological commitment issues

Imagine the sides have the following exchange:

Defense:	So you claim that the suspect is responsible for the damage?
Persecution:	Yes.
D:	And you agree that the claim that the sidewalk was unusually slippery that day due to an oil truck failure is at least as likely as that the suspect is responsible?
P:	Of course.
D:	And, given that you're accusing my client, you think you're in position to claim you know he is responsible, correct?
P:	Yes, that's correct.
D:	Are you not, then, in position to know that the sidewalk was unusually slippery that day due to an oil truck failure?
P:	...

I hope the reader shares the intuition that responding "no" to the above question would be a sign of a serious cognitive failure. In general, it seems that we'd normally expect the persecution to accept any claim relevant to the case that, given the evidence and the persecution's narration, is at least as likely as the guilt statement that they're putting forward. This motivates the following requirement:

(Commitment)   For any φ relevant to the case, if $\mathrm{P}^{f\mathbb{N}_i^A}(\varphi) \geq \mathrm{P}^{f\mathbb{N}_i^A}(\mathrm{G})$, then
$\mathrm{P}^{e\mathbb{N}_i}(N_i(\varphi)) \geq \mathrm{s}$.

If the reader prefers to do so, we might leave the notion of relevance at the intuitivel level. We don't have to though — it could be explicated along the following lines. First, a set of sentences is relevant for the case if it is consistent with the background knowledge and there is a narration such that its posterior probability given all background knowledge together with that set is different from its posterior probability given all background knowledge only. A set of sentences is a minimal relevant set if no proper subset thereof is a relevant set. A sentence is relevant if it or its negation is a member of a minimal relevant subset.

Now, how does (Commitment) apply to the gatecrasher? Take any g_i where $i \neq 1$. Consider an argument analogous to the previous one:

Defense:	So you claim that 1 is guilty?
Persecution:	Yes.
D:	And you agree 2 is at least as likely to be guilty as 1?
P:	Of course.
D:	So you're in position to claim you know 1 is guilty, correct?
P:	Yes, that's correct.
D:	Are you not, then, in position to know that 2 is guilty as well?
P:	...

Again, it seems that the negative answer to the above question would be irrational. (If you worry about g_i being relevant, wait for the development of the argument.)

Observe now that in the case of the gatecrasher, the accusing narration already discussed, $\mathbb{N}_1'$ is in the somewhat unusual lottery-paradox-like situation that given all that is known (and the narration itself) any other suspect is at least as likely to be guilty as suspect 1.

Now we can run the argument to the effect that N'_1 fails to satisfy conditions for conviction beyond reasonable doubt. Given the relevance of all g_i ($i \neq 1$), (Commitment) entails that for any $i \neq 1$ we have

$$P^{eN'_1}(N'_1(g_i)) \geq s. \tag{Step 1}$$

Since N'_1 is an accusing narration delivering g_1, by (Decision) and the fact that $a > s$ we have

$$P^{fN'_1}(g_i) \geq a > s. \tag{Step 2}$$

The very description of N'_1 entails:

$$g_i \notin N'_1. \tag{Step 3}$$

Steps (1-3) taken together, however, constitute the defining elements of (Gap) (for the slightly degenerate case where $\varphi_1 \vee \cdots \vee \varphi_u$ simply is g_i). This means that N'_1 is gappy, $\mathbf{G}(N'_1)$, and so it fails to justify a conviction beyond reasonable doubt.

What happens, on the other hand, if N'_1 is replaced with N^+_1: an accusing narration which results with taking the least narration containing N'_1 and closed under (Commitment) with respect to all g_i, $i \neq 1$, so that:

$$N^+_1 = \{g_1, e\} \cup \{g_k | k \neq 1\}? \tag{Bite the bullet}$$

Then, the resulting narration simply becomes highly implausible: e, and therefore also N^+_1, entails that exactly one person is innocent, but also for each particular suspect u, N^+_1 insists on u being guilty. Such a narration doesn't even satisfy (Initial plausibility), not to mention failing to be a dominating one (for which strong plausibility is required).

By the way, this is why each g_i is relevant: together they change the outcome. This is also the reason why they're relevant in the more technical sense: the set of all g_is contradicts evidence, while no proper subset thereof does.

So, to wrap up the discussion of the gatecrasher, in a sense there is no single simple diagnosis of the gatecrasher, for the following reasons. First of all, there is no single gatecrasher, but two main variants thereof: a realistic one and a very abstract one.[7] Second, some supposed solutions don't work against any of them. Third, sometimes there is more than one reason why a variant fails. Four, different variants can be seen as failing for different reasons (although reasons that apply to the abstract one apply to the realistic one; it's just that deploying them to the realistic gatecrasher is an overkill).

References

[1] Ronald J Allen (2010): *No Plausible Alternative to a Plausible Story of Guilt as the Rule of Decision in Criminal Cases*. In Larry Laudan Juan Cruz, editor: *Proof and standards of proof in the law*, Northwestern University School of Law, pp. 10–27.

[2] Ronald J Allen & Michael S Pardo (2007): *The problematic value of mathematical models of evidence*. The *Journal of Legal Studies* 36(1), pp. 107–140, doi:10.1086/508269.

[3] Jacob Bernoulli (1713): *Ars conjectandi*.

[4] Jonathan L. Cohen (1977): *The probable and the provable*. Oxford University Press, doi:10.2307/2219193.

---

[7]To be honest, multiple variants thereof depending on how unrealistic we're required to be. But let's stay content with two extreme cases.

[5] L Jonathan Cohen (1981): *Subjective Probability and the Paradox of the Gatecrasher*. Arizona State Law Journal, pp. 627–634.

[6] Mary Dant (1988): *Gambling on the truth: the use of purely statistical evidence as a basis for civil liability*. Columbia Journal of Law and Social Problems 22, pp. 31–70.

[7] Marcello Di Bello (2013): *Statistics and Probability in Criminal Trials*. Ph.D. thesis, University of Stanford.

[8] Susan Haack (2014): *Evidence Matters: Science, Proof, and Truth in the Law*. Cambridge University Press, doi:10.1017/CBO9781139626866.

[9] Susan Haack (2014): *Legal probabilism: an epistemological dissent*. In: *[8]*, Cambridge University Press, pp. 47–77.

[10] Hock Lai Ho (2008): *A philosophy of evidence law: Justice in the search for truth*. Oxford University Press, doi:10.1093/acprof:oso/9780199228300.001.0001

[11] David Kaye (1979): *The Paradox of the Gatecrasher and Other Stories*. Arizona State Law Journal, pp. 101–110.

[12] François Lepage (2012): *Partial Probability Functions and Intuitionistic Logic*. Bulletin of the Section of Logic 41(3/4), pp. 173–184.

[13] François Lepage & Charles Morgan (2003): *Probabilistic Canonical Models for Partial Logics*. Notre Dame Journal of Formal Logic 44(3), pp. 125–138, doi:10.1305/ndjfl/1091030851.

[14] Charles R. Nesson (1979): *Reasonable doubt and permissive inferences: The value of complexity*. Harvard Law Review 92(6), pp. 1187–1225, doi:10.2307/1340444.

[15] Nancy Pennington & Reid Hastie (1992): *Explaining the evidence: Tests of the Story Model for juror decision making*. Journal of personality and social psychology 62(2), pp. 189–204, doi:10.1037/0022-3514.62.2.189.

[16] Nancy Pennington & Ried Hastie (1991): *A cognitive theory of juror decision making: The story model*. Cardozo Law Review 13, pp. 519–557.

[17] Alex Stein (2005): *Foundations of Evidence Law*. Oxford University Press, doi:10.1093/acprof:oso/9780198257363.001.0001.

[18] Laurence H Tribe (1971): *A Further Critique of Mathematical Proof*. Harvard Law Review 84, pp. 1810–1820, doi:10.2307/1339570.

[19] Laurence H Tribe (1971): *Trial by mathematics: Precision and ritual in the legal process*. Harvard Law Review 84(6), pp. 1329–1393, doi:10.2307/1339610.

[20] Barbara D. Underwood (1977): *The thumb on the scale of justice: Burdens of persuasion in criminal cases*. Yale Law Journal 86(7), pp. 1299–1348, doi:10.2307/795788.

[21] Gary L Wells (1992): *Naked statistical evidence of liability: Is subjective probability enough?* Journal of Personality and Social Psychology 62(5), pp. 739–752, doi:10.1037/0022-3514.62.5.739.