Quaestio facti. Revista Internacional sobre Razonamiento Probatorio
Quaestio facti. International Journal on Evidential Legal Reasoning
Sección: Conjeturas y refutaciones
2026 | 11 pp. 157-180
Madrid, 2026
DOI: 10.33115/udg_bib/qf.i11.23278
Marcial Pons Ediciones Jurídicas y Sociales
© Marcello Di Bello
ISSN: 2604-6202
Recibido 5/5/26 | Aceptado 25/5/26 | Publicado *online*: __/__/2026
Editado bajo licencia Reconocimiento 4.0 Internacional de Creative Commons

EVENTS AND IDENTITY IN PROBABILISTIC MODELS OF LEGAL EVIDENCE

Marcello Di Bello
Arizona State University

ABSTRACT: This paper examines whether Bayesian networks are expressive enough to model reasoning with evidence in legal cases. Bayesian networks can represent many familiar patterns of evidential reasoning, including inferences from evidence to hypotheses, cumulative support from multiple items of evidence, and chains of inferences linking intermediate hypotheses to ultimate guilt. But other evidential inferences common in legal cases are more difficult to model. Focusing on a real criminal case, the paper distinguishes between identity- and event-level inferences. Event-level inferences show why certain actions amount to guilty conduct, while identity-level inferences link the defendant to those actions. The challenge for Bayesian network models of legal evidence is to represent how identity- and event-level inferences combine and reinforce one another. Meeting this challenge requires extending Bayesian networks beyond a purely propositional language.

KEYWORDS: bayesian networks; likelihood ratio; probability; criminal law.

1. INTRODUCTION

Can the evidence in a legal case be modeled probabilistically? Following Mackor (2026), this question is best divided into more specific subquestions. One concerns the target of the probabilistic modeling: are we asking whether individual items of evidence can be modeled probabilistically, or is the question about modeling the totality of the evidence in a case? Another subquestion concerns who is supposed to carry out the probabilistic analysis. Should the analysis be assigned to experts who possess domain-specific knowledge, or should it remain with judges and jurors, perhaps with the assistance of court-appointed experts?

There is little doubt that individual items of evidence can be—and often are—evaluated probabilistically by experts. The most familiar examples come from forensic science (Taroni *et al.*, 2014). Suppose that genetic material matching the defendant is found at the crime scene. A forensic expert will typically offer the following probabilistic analysis of the match: they will assess how probable that finding would be if the defendant were the source of the genetic material, as compared with how probable it would be if someone else were the source. Expressed as a ratio, this comparison is known as the likelihood ratio:

$$\frac{\boldsymbol{P}\,(\textbf{ match} \mid \textbf{defendant is the source})}{\boldsymbol{P}\,(\textbf{ match} \mid \textbf{defendant is not the source})}$$

This ratio quantifies the degree of support that the match evidence provides for the source hypothesis: the greater the ratio, the stronger the support.

Forensic experts have also begun to apply probabilistic analyses beyond source-level propositions to activity-level propositions (Taylor *et al.*, 2018). The question of interest is often not merely whether the defendant is the source of a trace recovered at the crime scene, but how the trace came to be there, for instance, whether it got there during a violent confrontation between the defendant and the victim. How to carry out a probabilistic analysis of trace evidence for activity-level propositions is a topic of ongoing discussion in the literature (Stacey *et al.*, 2025).

A typical legal case, however, will consist of many pieces of evidence bearing on different propositions, not a single item of evidence bearing on a single proposition, such as a source- or activity-level hypothesis. So it is natural to ask whether the totality of the evidence in a legal case can be modeled probabilistically. That is a more complicated question. Which also makes the institutional question—who should carry out the probabilistic analysis?—more complicated. Experts may have the technical competence to build probabilistic models of individual items of evidence, but they lack the institutional authority to model a legal case as a whole. Judges and jurors, by contrast, have the authority to assess the totality of the evidence—that is what they are tasked with doing at trial—but they usually lack the technical competence required to construct a probabilistic model.

In this paper, I am going to set aside the institutional question. I will focus instead on whether it is in principle possible to model a case as a whole in probabilistic terms. A number of probabilistic models of entire legal cases have been developed. These models all rely on Bayesian networks (more on these soon). Examples include analyses of the Sacco and Vanzetti case (Kadane & Schum, 1996), the Anjum murders (Vlek *et al.*, 2014), the Simonshaven case (Fenton *et al.*, 2020), and a recent supermarket robbery case in the Netherlands (Hampson & Leeuwen, 2025). These examples suggest that probabilistic methods can be extended beyond individual items of evidence to the analysis of a case as a whole using Bayesian networks. So there are some reasons for optimism.

But there are also reasons for skepticism. First, a common worry is that probabilistic models require numerical probabilities as inputs, yet the relevant data are often unavailable. Without data, the analysis must rely on rough estimates, educated guesses or expert judgment. But even when the data are available, they need not dictate the relevant probabilities. The data must first be interpreted and connected to the propositions at issue in the case. As Allen & Pardo (2007) note, the numbers used in probabilistic models are not "an objective datum that captures the value of evidence". Instead, they are "just more evidence, which itself must be interpreted to assess its value" (p. 119) [1].

A second challenge for probabilistic analyses of evidence, especially when the aim is to model an entire case, is the absence of a well-established methodology. Different analysts may produce different probabilistic models of the same case. To be sure, some authors have proposed ways to regiment the process of probabilistic model construction, for example, by identifying recurring idioms (Fenton *et al.*, 2013) or by comparing the modeling choices of different analysts (Hampson & Leeuwen, 2025). But the construction remains an interpretive task rather than a mechanical procedure. Any model—and probabilistic models are no exception—reflects choices about which propositions to include, at what level of abstraction to represent them, and how to structure the relations among them.

Now, even if the numerical and methodological challenges could be addressed, a third challenge would remain: are probabilistic models—and Bayesian networks in particular—expressive enough to capture the kinds of inferences that characterize reasoning with evidence in legal cases? This challenge is my focus here. As I will show, some evidential inferences can be represented quite naturally in Bayesian networks, while others are less tractable. To anticipate, the difficulty I will highlight lies in the interaction between inferences about events and inferences about identity. The integration of the two is not easy to model within standard (propositional) Bayesian networks. A richer formalism is needed.

The plan is as follows. I begin with a few examples of inferences that can be modeled with Bayesian networks (§2 and §3). To keep the discussion concrete, I turn to a simple legal case. I start with an informal description of the case and its items of

[1] A subjectivist Bayesian will respond that this is not a defect. Probabilities are not objective frequencies mechanically read off from the data. They are instead subjective degrees of belief, informed by data where available and revised in light of new evidence. On this view, probabilistic models do not remove judgment from the assessment of the evidence. Rather, they make judgment explicit and require it to be probabilistically coherent (Bozza *et al.*, 2026). Still, the probabilistic model does not, by itself, tell us whether the probability assignments it contains are well supported by the evidence. That further question depends on how the evidence is interpreted, how the propositions are framed, and whether the available data are relevant to the case at hand.

evidence (§4). I then present a coarse probabilistic model and refine it (§5 and §6). I show how some inferences resist probabilistic modeling in standard (propositional) Bayesian networks (§7 and §8). Finally, I conclude by underscoring the need for a more expressive formalism (§9).

2. BAYESIAN NETWORKS

I shall start with an overview of the central modeling tool used in probabilistic analyses of legal evidence: Bayesian networks. They are graphical representations of probabilistic dependencies among a set of variables. A Bayesian network has both a qualitative, graphical component and a numerical component[2]. Graphically, the network consists of nodes and directed edges (arrows) connecting the nodes. The graph must be acyclic (Figure 1): one cannot start from a node and, by following the direction of the arrows, eventually return to that same node.

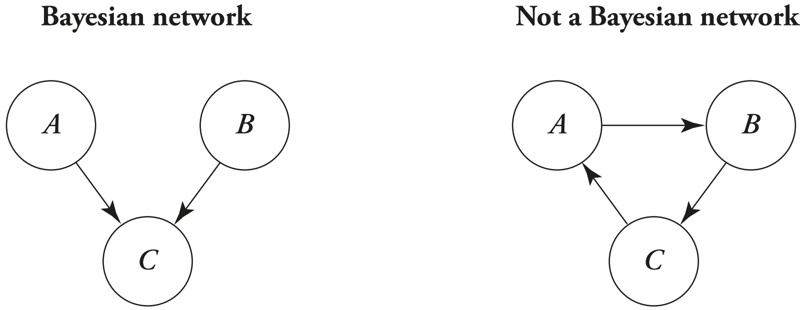

Figure 1. A Bayesian network must be directed and acyclic. The graph on the left satisfies both conditions. The graph on the right is directed, but not acyclic.

On the numerical side, each node in a Bayesian network represents a random variable that can take different possible values with different probabilities. Each node is associated with a probability table that specifies the probabilities of its possible values (Figure 2). In other words, the table specifies the probability distribution of the random variable. The directed edges indicate conditional probabilistic dependencies among the variables. If a node has no parents, its probability table will give a prior probability distribution. If a node has parents, its probability table will give a conditional probability distribution, specifying the probability of each possible value of that node given the values of its parents.

[2] For a book length introduction to Bayesian networks for modeling legal cases, see (Lagnado, 2021). For the application of Bayesian networks to forensic science, see (Taroni *et al.*, 2014).

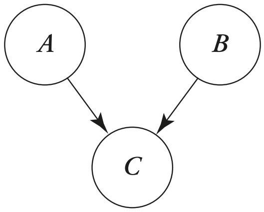

$P(A = 1) = 0.3$, $P(A = 1) = 0.7$

$P(B = 1) = 0.4$, $P(B = 1) = 0.6$

A	B	$P(C = 1 \mid A,B)$	$P(C = 0 \mid A,B)$
0	0	0.10	0.90
0	1	0.50	0.50
1	0	0.60	0.40
1	1	0.90	0.10

Figure 2. Nodes ***A*** and ***B*** have no parents and are assigned prior probabilities. Node ***C*** has parents ***A*** and ***B***, so it is assigned a conditional probability table.

A common source of confusion concerns the direction of the arrows. Strictly speaking, the difference between a network with an arrow from ***A*** to ***B*** and a network with an arrow from ***B*** to ***A*** is purely formal. In the first case, the model requires a conditional probability distribution for ***B*** given ***A***. In the second case, the model requires a conditional probability distribution for ***A*** given ***B***.

But arrows are often given an intuitive, less formal interpretation. On a common reading, an arrow from ***A*** to ***B*** suggests that ***A*** is a cause, or part of the causal explanation, of ***B***. For example, whether the defendant is the source causally influences whether the forensic test produces a match. So the graph will have an arrow from source to match. From an evidential point of view, however, the observed match is a reason for thinking that the defendant is the source. On that reading, the arrow will run from match to source (Dahlman & Kolflaath, 2022). In what follows, I will often use the causal reading. When evidence is modeled as the effect of an event, the arrow will run from the event to the evidence.

Arrows can also be given an analytical interpretation. For example, one node may be materially or conceptually incompatible with another: if one proposition is true, the other must be false. A node for the defendant's alibi and a node for the defendant's guilt are related in this way: if the defendant was elsewhere at the time of the offense, then the defendant could not have materially committed the offense. Or one node may be a consequence of another: if the defendant disposed of the weapon at a particular location, then the defendant must have been at that location. These relations reflect physical or conceptual constraints, and they can be represented in a Bayesian network by deterministic entries in the relevant conditional probability tables (Figure 3, bottom).

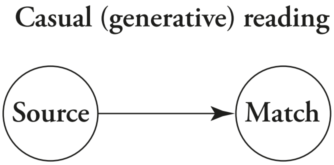

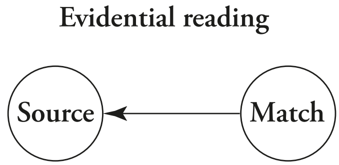

G	A	$P(C = 1 \mid G, A)$
1	1	0
1	0	1
0	1	1
0	0	1

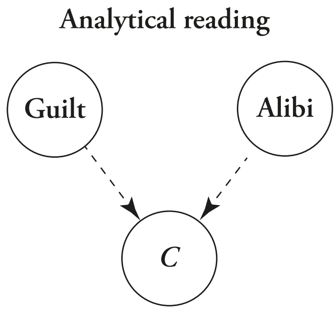

Figure 3. In a causal reading, the arrow runs from cause to effect. In an evidential reading, reasoning proceeds in the opposite direction. The bottom panel illustrates an analytical constraint (note the dashed arrows): an alibi and guilt are materially incompatible. This can be represented by a constraint node C, whose conditional probability table assigns probability zero to the combination of guilt and alibi.

Bayesian networks provide compact representations of joint probability distributions over many variables. The basic idea is simple: instead of writing out an enormous probability distribution that covers every possible combination of variables, the network breaks the problem into smaller pieces. Each variable is represented by a node, and the probability of that variable depends only on its parent variables, that is, the variables that point directly into it. Stated more formally, the joint probability distribution over all variables in the network can be factorized as a product of local conditional probabilities, each associated with a node and its parents. If the variables are $X_1, \ldots, X_n$, the joint distribution is written as the following product:

$$P(X_1, \ldots, X_n) = \prod_{i=1}^{n} P(X_i \mid \mathbf{Pa}(X_i))$$

where $\mathbf{Pa}(X_i)$ denotes the parents of X_i in the graph. The formula expresses the central assumption of a Bayesian network: once the parents of a variable are fixed, other non-descendant variables do not make any further difference to the probability of that variable. To put it more formally, the equality holds by the local Markov property of Bayesian networks: each node is conditionally independent of its non-descendants given its parents (Figure 4).

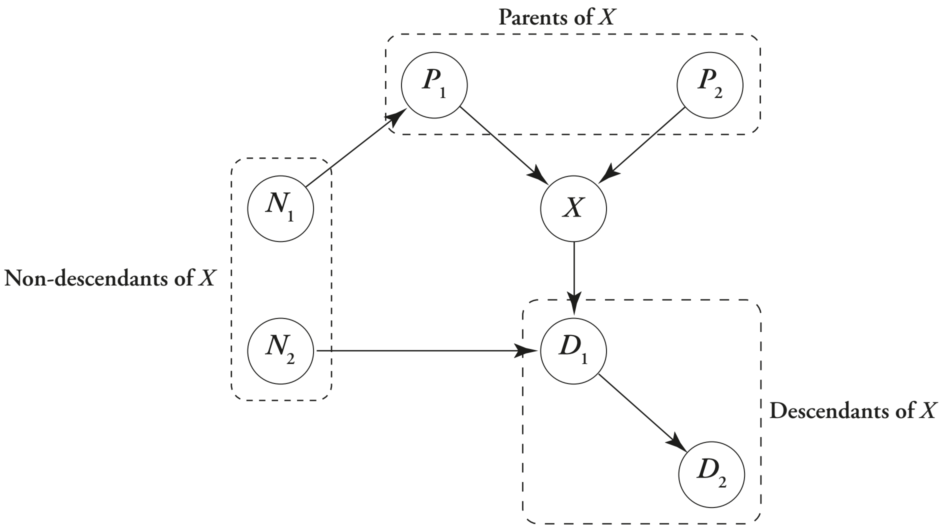

Figure 4. Once X's parents P_1 and P_2 are fixed, X's non-descendants N_1 and N_2 provide no further information about X. So, by the local Markov property, it follows that $P(X \mid P_1, P_2, N_1, N_2) = P(X \mid P_1, P_2)$.

3. IDIOMS IN BAYESIAN NETWORKS

It is now time to consider how Bayesian networks can model recurring patterns of inference in legal cases. In the diagrams that follow, some nodes are evidence nodes (represented by rectangles) while others are hypothesis nodes (represented by circles). A Bayesian network can thus be understood as a structure for representing probabilistic relationships or inferences between evidence and hypotheses.

In what follows, I will slightly abuse notation. I will use the same symbol both for the random variable represented by a node and for the proposition that this variable takes the value **1**. Thus, if S is a node representing the source hypothesis, then S denotes the random variable that can take the values $S = 1$ and $S = 0$. At the same time, I will use S as shorthand for the proposition that the source hypothesis is true (that is, $S = 1$) and $\neg S$ as shorthand for the proposition that it is false (that is, $S = 0$). Context should make clear which use is intended.

Following recurring inference patterns or idioms for Bayesian networks (Fenton *et al.*, 2013), I will start with the simplest one. The support that a piece of evidence lends to a hypothesis can be modeled by comparing the probability of observing the evidence if the hypothesis is true with the probability of observing the same evidence if the hypothesis is false. As seen before, this is the likelihood ratio:

$$\frac{P(E \mid H)}{P(E \mid \neg H)}$$

If the ratio is greater than one, a rational decision-maker should increase their assessment of the probability of the hypothesis[3]. Graphically (figure 5, left side), this inference is modeled by the simplest Bayesian network consisting of just two nodes, one for the hypothesis and another for the evidence, with an arrow from hypothesis to evidence. Crucially, an arrow from H to E does not by itself mean that E supports H. That depends on the corresponding conditional probability table. If the table says that $P(E \mid H) > P(E \mid \neg H)$, then E supports H; if $P(E \mid H) = P(E \mid \neg H)$, then E is evidentially neutral; and if $P(E \mid H) < P(E \mid \neg H)$, then E counts against H. Thus, the graph signals a potential dependency. Whether that dependency amounts to support, neutrality, or opposition depends on the numbers in the probability table.

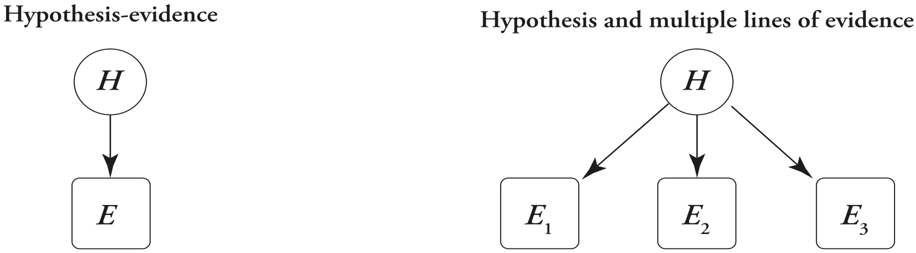

Figure 5. Simple inference patterns in Bayesian network form.

A more complex pattern arises when multiple items of evidence, $E_1, E_2, \ldots, E_k$, constitute independent lines of support for the same hypothesis H. This evidential pattern has a formal equivalent, that is, each E_i can be assumed to be conditionally independent given the hypothesis of interest (and given its negation)[4]. Then, their combined evidential force is modeled by multiplying their likelihood ratios:

[3] By the odds form of Bayes' theorem,

$$\frac{P(H \mid E)}{P(\neg H \mid E)} = \frac{P(E \mid H)}{P(E \mid \neg H)} \cdot \frac{P(H)}{P(\neg H)}$$

If the likelihood ratio is greater than one, then the posterior odds of H are greater than the prior odds of H. Provided $0 < P(H) < 1$, it follows that $P(H \mid E) > P(H)$, or in other words, the observation of the evidence rationally supports an increased degree of belief in the hypothesis.

[4] In symbols, for three items of evidence E_1, E_2, E_3:
$P(E_1,E_2,E_3|H) = P(E_1|H)P(E_2|H)P(E_3|H)$,
$P(E_1,E_2,E_3|\neg H) = P(E_1|\neg H)P(E_2|\neg H)P(E_3|\neg H)$

$$\frac{P(E_1|H)}{P(E_1|\neg H)} \times \frac{P(E_2|H)}{P(E_2|\neg H)} \times \ldots \times \frac{P(E_k|H)}{P(E_k|\neg H)}$$

If each likelihood ratio is greater than one, a rational decision-maker should increase their assessment of the probability of the hypothesis by a significant margin. Multiple, independent lines of evidence in support of ***H*** lend much stronger support taken together than each one in isolation. Graphically (figure 5, right side), this inference is modeled by a Bayesian network consisting of a hypothesis node from which outgoing arrows point to different evidence nodes.

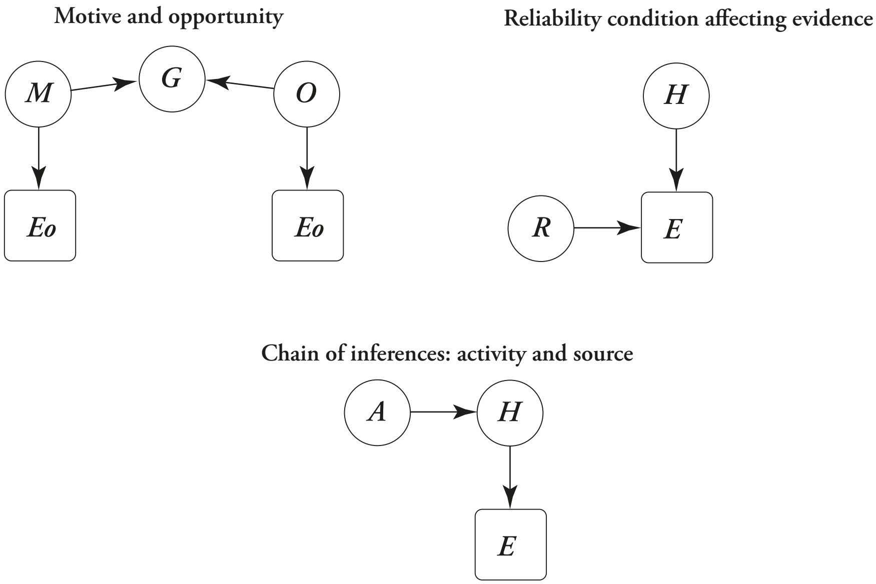

Figure 6. Other inference patterns in Bayesian network form.

Fenton *et al.* (2013) identify many other recurring inference patterns, including inferences from someone's motive and opportunity to their guilt. Having a motive and having the opportunity to commit a crime, for example by being near the crime scene, are both predictors of guilt. In a Bayesian network, this can be represented by nodes for motive and opportunity, each with an arrow pointing to the guilt node. Motive and opportunity may in turn be supported by evidence, so that they function as intermediate hypotheses within the network (Figure 6, top left side).

Another common pattern is the reliability (or accuracy) idiom. When a piece of evidence supports a hypothesis, its probative value depends on the reliability of the evidence. Suppose, for example, that an eyewitness states that the defendant was near the crime scene. That testimony supports the hypothesis that the defendant was near the scene only if the witness is reliable, say was not far away or did not make the identification in the dark. Graphically (figure 6, top right side), this can be represented by a network in which both the evidence node and a reliability node bear on the hypothesis. In likelihood ratio terms:

$$\frac{P(E \mid H,R)}{P(E \mid \neg H,R)} > 1 \quad \frac{P(E \mid H,\neg R)}{P(E \mid \neg H,\neg R)} = 1$$

The first expression says that, conditional on the witness being reliable, the testimony is more probable if the hypothesis is true than if it is false. So the evidence supports the hypothesis. The second says that, conditional on the witness being unreliable, the testimony is equally probable whether the hypothesis is true or false. So the evidence provides no support when the reliability condition is not met.

Finally, consider the inference from match evidence to a source proposition, and from the source proposition to an activity proposition (Figure 6, bottom). The inference proceeds in two steps. First, the match evidence supports, to some degree, the source proposition, for example, that the trace recovered at the crime scene came from the defendant. Second, the source proposition supports, to some degree, an activity proposition, for example, that the defendant participated in the relevant criminal conduct. These two inferential steps can come apart. A DNA match could strongly support the hypothesis that the defendant was the source of the trace, while it remains unclear whether the trace was deposited during the crime or through innocent contact. Conversely, even if the source-to-activity inference is strong, the overall inference can still be weak if the matching evidence provides only limited support for the source proposition.

The inferential patterns identified above—the simple evidence-hypothesis idiom; multiple lines of evidence; motive and opportunity; reliability; chain of inferences—can also be combined in a modular fashion to form more complex models that consist of several items of evidence and hypotheses. Hampson & Leeuwen (2025) rely on those very idioms to model a supermarket robbery, and Fenton *et al.* (2020) do something similar with the Simonshaven case. This suggests that Bayesian networks have considerable expressive power: by combining common inferential idioms, they can model some legal cases in their entirety.

4. STATE VERSUS EVANS

Having seen some of the inferences that Bayesian networks can model, let us now turn to an actual legal case to keep the discussion concrete. The case is useful for our discussion because it is simple to understand but also complex enough to pose a challenge for probabilistic modeling. The prosecution's theory depends on multiple strands of evidence that support different propositions. Some items bear on motive, some on identity, some on presence near the crime scene, and others on possible post-offense conduct.

In *State v. Evans*, SC 21006 (Conn. Aug. 12, 2025), the prosecution alleged that the defendant, Evans, shot and killed Reginald May outside May's apartment in Bridgeport, Connecticut, shortly after 4:00 a.m. on July 2, 2017. The State's theory was that the killing was connected to an earlier dispute between Evans and the victim. Several items of evidence, taken together, link Evans to the shooting.

Evans and May had argued and fought over money during a moving job in Virginia on June 28, 2017. After that incident, May returned to Connecticut. A few days later, surveillance video from near May's apartment showed a large dark SUV circling the area and then parking on Alice Street at about 4:17 a.m. The footage showed a man in a hooded sweatshirt exiting the SUV, walking toward the apartment building, and returning to the vehicle at about 4:21 a.m.

The State sought to connect this footage to Evans in several ways. First, the SUV shown in the video was consistent with Evans's Cadillac Escalade. Second, police extracted a still image from the surveillance footage, and both the victim's brother, John May, and the victim's girlfriend, Cherry Williams, identified the person in the still image as Evans. The State also introduced evidence of a firearm recovered on July 13, 2017, at West Rock Nature Center. Ballistic examination linked the firearm to the crime. Finally, the State relied on phone location information showing that on July 3, 2017, Evans's phone connected to a cell tower near West Rock Ridge State Park, near the area where the firearm linked to the crime was found.

A list of the different items of evidence is provided below for ease of reference:

Surveillance video. Surveillance footage on July 2, 2017 from near the victim's apartment showed a dark SUV parking on Alice Street around 4:17 a.m. A man in a hooded sweatshirt exited the SUV, walked toward the building, and returned to the SUV around 4:21 a.m.

ID from video still. A still image extracted from the surveillance footage was shown to John May and Cherry Williams, both of whom identified the person in the image as Evans.

SUV evidence. SUV in the video was consistent with Evans's Cadillac Escalade.

Firearm evidence. A firearm and extended magazine were recovered on July 13, 2017, at West Rock Nature Center, followed by ballistic examination linking the firearm to the shooting.

Phone location. On July 3, 2017, Evans's phone connected to a cell tower near West Rock Ridge State Park, close to the location where the firearm connected to the crime was recovered.

5. INDEPENDENT LINES OF EVIDENCE

As a first attempt, we can model this case with a Bayesian network that uses the five key pieces of evidence all at once: ID, SUV, firearm, phone and surveillance video evidence. Each of them links the defendant to the crime in its own way: via the phone location and recovery of the firearm in that location; similarity of the SUV in the surveillance video with defendant's vehicle; and identification of the person in the video. The structure in Figure 7 illustrates the textbook case for aggregating independent items of evidence, all bearing on the same hypothesis.

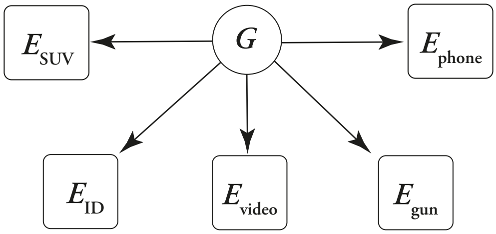

Figure 7. Simplest Bayesian network: ID, SUV, phone, firearm and video evidence are treated as conditionally independent indicators bearing on guilt.

Quantitatively, each piece of evidence can be paired with a measure of how strongly, on its own, it supports the guilt hypothesis, such as the likelihood ratio:

$$LR(E_i;G)=\frac{P(E_i \mid G)}{P(E_i \mid \neg G)}$$

In this model, the combined strength of support of the five pieces of evidence in favor of the guilt hypothesis is the product of the individual likelihood ratios:

$$LR(E_{ID};G)\cdot LR(E_{SUV};G)\cdot LR(E_{PHONE};G)\cdot LR(E_{GUN};G)\cdot LR(E_{VIDEO};G)$$

Suppose that the ID, SUV, phone, firearm and video evidence have likelihood ratios of 3, 5, 4, 2 and 2. The numbers are illustrative. Then, their combined likelihood ratio will be **3×5×4×2×2** = **240**. To have a sense of the strength of such evidence, in the aggregate, a likelihood ratio of 240 would be able to turn a 1% initial probability of guilt (1-to-99 prior odds) into approximately 70% posterior probability of guilt (2.4-to-1 posterior odds)[5].

But there is a complication. Our model treats the evidence as a collection of independent items all bearing on the guilt hypothesis, as if each were a standalone indicator of guilt. In fact, the different pieces of evidence speak to different hypotheses. The phone evidence links the defendant to the crime by showing that the perpetrator disposed of the weapon in the same area that the defendant visited. The video, ID and SUV evidence place the defendant outside a building at a certain time. The video

[5] The calculation uses the odds form of Bayes' theorem: posterior odds = prior odds × likelihood ratio. A prior probability of **1%** corresponds to prior odds **1/99**. Multiplying by **240** gives posterior odds **240/99≈2.42**. Converting odds back into probability gives **2.42/(1+2.42)≈71%**.

shows what the person outside the building did; the ID and SUV evidence connect the person in the video to the defendant. Thus, the model of multiple independent lines of evidence, each supporting guilt, is too simplistic.

6. OPPORTUNITY → GUILT → DISPOSAL

To refine our model, I shall now distinguish between the opportunity hypothesis and the guilt hypothesis, and between the guilt hypothesis and disposal hypothesis. The opportunity hypothesis (OPP) is the claim that Evans was present near the crime scene around the time of the shooting. The guilt hypothesis (G) is the stronger claim that Evans committed the shooting. The surveillance video ID and the SUV evidence are naturally interpreted as evidence that Evans was in the relevant location at the relevant time and thus had the opportunity to commit the crime. Evans' guilt is supported through the further inference from OPP to G. We can represent this structure as a Bayesian network (Figure 8, right side) in which both evidentiary nodes support OPP, and OPP in turn supports G.

Next, we will add two other evidence nodes: the recovery of the firearm and the post crime phone location evidence. The two items of evidence support—how, exactly, needs to be seen—the intermediate proposition that Evans disposed of the weapon. So we can add a disposal node to the network, call it ***D***. We can complete the model by drawing arrows from the disposal node to the evidence nodes, one for evidence about where the weapon was found and another for the phone location. The disposal hypothesis supports guilt via further inference.

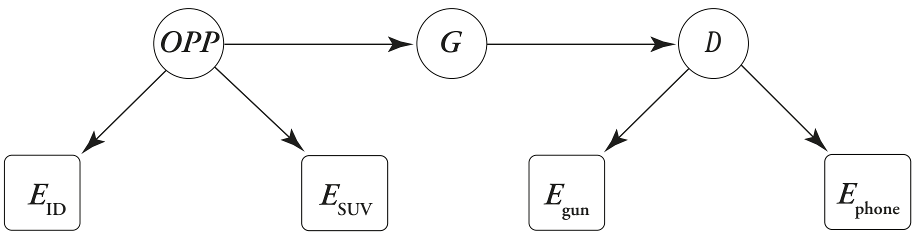

Figure 8. Bayesian network distinguishing guilt (***G***) from post-crime disposal of the firearm (***D***) and opportunity (***OPP***). The evidence of the recovered firearm and the phone-location evidence most naturally bear on disposal, which in turn supports guilt. The ID and SUV evidence, instead, most naturally bear on opportunity.

The new model (Figure 8) is more faithful to the evidential structure of the case than a flatter model in which each evidential item supports guilt. The new model makes explicit that the evidence does not bear on guilt in a simple one-step way. Instead, the evidence supports intermediate events, and these in turn support guilt. We have two independent evidential routes, one that goes through opportunity, and another that goes through the disposal hypothesis. The contributions of each route add up. But how, exactly? The math here is more involved than before.

Let us start with the opportunity route. In the model, there are two independent lines of evidence bearing on the opportunity hypothesis, E_{ID} and E_{SUV}. (I'll revisit this modeling assumption later.) By conditional independence, the likelihood ratio for the opportunity hypothesis is

$$LR(E_{ID}, E_{SUV};OPP)= \frac{P(E_{ID}|OPP)}{P(E_{ID}|\neg OPP)} \cdot \frac{P(E_{SUV}|OPP)}{P(E_{SVU}|\neg OPP)}$$

The problem now is that we need to make a further inference from opportunity to guilt. The likelihood ratio relative to opportunity, based on the ID and SUV evidence, must be transferred to the guilt hypothesis. This is the formula we need[6]:

$$LR(E_{ID}, E_{SUV};OPP)= \frac{LR(E_{ID}, E_{SUV};OPP)\ P(OPP|\ G)+1\ P(OPP|\ G)}{LR(E_{ID}, E_{SUV};OPP)\ P(OPP|\ \neg G)+1\ P(OPP|\ \neg G)}$$

We can run a similar argument for the evidence supporting the disposal hypothesis which in turn supports the guilt hypothesis. This is the other evidential route. In the model, there are two items of evidence bearing on the disposal hypothesis, E_{gun} and E_{phone}. Again by conditional independence, the combined likelihood ratio for the disposal hypothesis is

$$LR(E_{\mathrm{gun}}, E_{\mathrm{phone}};D)= \frac{P(E_{\mathrm{gun}}|D)}{P(E_{\mathrm{gun}}|\neg D)} \cdot \frac{P(E_{\mathrm{phone}}|D)}{P(E_{\mathrm{phone}}|\neg D)}$$

As before, the likelihood ratio relative to the disposal hypothesis, based on gun and phone evidence, must also be transferred to the guilt hypothesis:

$$LR(E_{\mathrm{gun}}, E_{\mathrm{phone}};G)= \frac{LR(E_{\mathrm{gun}}, E_{\mathrm{phone}};D)\ P(D|G)+1 - P(D|G)}{LR(E_{\mathrm{gun}}, E_{\mathrm{phone}};D)\ P(D|\neg G)+1 - P(D|\neg G)}$$

So, evidence that more closely supports ***OPP*** must be translated into support for ***G*** through the opportunity-to-guilt route; evidence that more closely supports ***D*** must also be translated into support for ***G*** through the disposal-to-guilt route. As a rule of thumb, the inference toward guilt along each route is constrained by

---

[6] Let $E = (E_{ID}, E_{SUV})$.Assume that ***E*** bears on guilt ***G*** only through the opportunity proposition ***OPP***. Then, by the law of total probability,

$$\frac{P(E\,|G)}{P(E\,|\neg G)} = \frac{P(E\,|\,OPP)\ P(OPP\,|\,G) + P(E\,|\,\neg OPP)\ P(\neg OPP\,|\,G)}{P(E\,|\,OPP)\ P(OPP\,|\,\neg G) + P(E\,|\,\neg OPP)\ P(\neg OPP\,|\,\neg G)}$$

Dividing numerator and denominator by $P(E\,|\,\neg OPP)$ we obtain the formula.

the weakest inferential link. In other words, the evidential force transmitted to ***G*** is bounded both by how strongly the evidence supports ***OPP*** and by how strongly ***OPP*** supports ***G***. A weak link at either stage acts as a bottleneck. More formally,

$$LR(E_{ID}, E_{SUV}; G) \leq \min \left\{ LR(E_{ID}, E_{SUV}; OPP), \frac{P(OPP \mid G)}{P(OPP \mid \neg G} \right\}$$

Similarly, along the route from disposal to guilt, the inference is constrained by the weakest inferential link:

$$LR(E_{gun}, E_{phone}; G) \leq \min \left\{ LR(E_{gun}, E_{phone}; D), \frac{P(D \mid G)}{P(D \mid \neg G} \right\}$$

Finally, the overall likelihood ratio for the guilt hypothesis, based on all the evidence we have, is the result of the following multiplication:

$$LR(E_{ID}, E_{SUV,} E_{gun}, E_{phone}; G) = LR(E_{ID}, E_{SUV,}; G) \cdot LR(E_{gun}, E_{phone}; G)$$

This holds because the items of evidence along each evidential route are conditionally independent given the guilt hypothesis, or in other words, the opportunity-related evidence and disposal-related evidence each provide an independent route of support for guilt. So the support for ***G*** adds up across the two evidential routes and the relevant likelihood ratios multiply. But since each route is bounded by the weakest inferential link, when the two routes are aggregated, the overall strength of the inference is bounded by the product of the weakest links in each route. All in all, the total evidential force accumulate across independent evidential routes, but only after each route has been limited by its weakest inferential link.

7. EVENT AND IDENTIFICATION INFERENCES

The model in the previous section, albeit an improvement on the earlier one, suffers from a number of problems. It incorporates four sources of evidence: SUV, ID, phone location and firearm evidence. But it leaves out a fifth, namely the content of the surveillance video. The video shows a dark SUV arriving near the apartment complex shortly before the shooting, a person exiting the SUV and walking toward the apartment complex, and that person returning shortly afterward. From the sequence, location and timing of events, we can infer that the person who engaged in the actions depicted in the video almost certainly committed the shooting. But, from a modeling standpoint, it is not clear where the video evidence should be placed in the inferential structure represented by the network. It does not bear in any natural way on the opportunity hypothesis, at least not if opportunity is understood as the proposition that Evans was near the crime scene. The video evidence, by itself, does not link Evans to the crime. It only links the person in the video to the shooter, but leaves open who that person may be.

These observations also suggests that the model misrepresented the inferential role of the ID and SUV evidence. The model takes those items to be evidence for the opportunity hypothesis that Evans was near the crime scene. But, ID and SUV evidence do not independently support that hypothesis. Rather, they help identify the person in the video as Evans and the vehicle as his SUV. Only when that identification is combined with the content of the video can we infer that Evans was near the crime scene at the relevant time.

A third problem—related to the second—is that the proposed model misleadingly suggests that phone and firearm evidence, separately, support the disposal hypothesis. In fact, one without the other would be irrelevant. The recovery of the firearm at the remote location does not by itself show who placed it there. Nor does the suspect's phone being located near that site, by itself, establish involvement in the disposal of the firearm. Each item seems to derive much of its significance from the other. The firearm evidence makes the phone evidence probative because the suspect appears near the place where the crime weapon was found. The phone evidence makes the firearm evidence probative because the suspect appears at the place where the weapon was apparently discarded. The force of the evidence therefore arise from their fit within a narrative of post-crime disposal.

One way to diagnose what is going on is to distinguish two distinct inferential tasks, and possibly more. The first task is event reasoning: what happened, how, where and when? Here the aim is to determine the sequence of events in space and time. The video and firearm evidence address event reasoning. Someone approached the location, as shown in the video; the victim was killed; the firearm used in the crime was later disposed of in a remote location; and the recovered firearm is linked to the crime by forensic analysis. The second task is identity reasoning: who did it? Once the event sequence is fixed, the question becomes whether the defendant is the person who carried out those actions. The SUV and ID evidence address the identity question prior to the crime: who was the person around the crime scene? The phone location evidence, instead, addresses the identity question post-crime, specifically whether the defendant was the person who went to the disposal site[7].

If this analysis is right, a more adequate model should clearly distinguish hypotheses about events from hypotheses about identity, and then combine the two. One way to proceed is to build two preliminary submodels (Figure 9). The event model will contain propositions such as: an individual approached the scene, killed the victim, and later disposed of the firearm at the remote location. The identity model would then ask whether that individual was the defendant. The central challenge is then to explain how these two submodels can be combined.

[7] Tribe (1971) distinguishes between (1) evidence "directed to the occurrence or nonoccurrence of the event, act, or type of conduct on which the litigation is premised" and (2) evidence "directed to the identity of the individual responsible for a certain act or set of acts." (p. 1339)

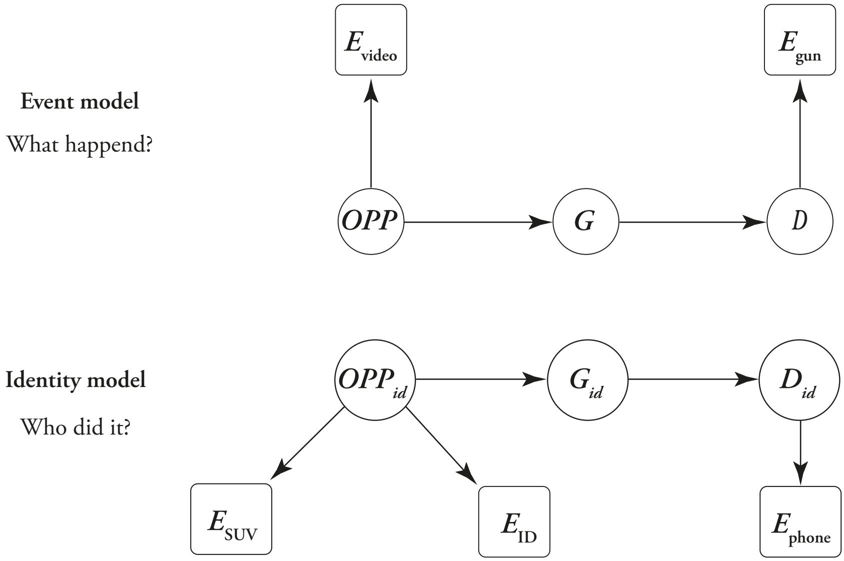

Figure 9. Two Bayesian networks with the same structure, $\boldsymbol{OPP \rightarrow G \rightarrow D}$, but different interpretations. In the event model, the nodes represent propositions about the sequence of events: opportunity, commission of the crime, and later disposal of the firearm. In the identity model, the same nodes represent propositions about whether the defendant is the relevant actor at each stage.

8. COMBINING THE TWO

In an attempt to build a unified model that does not confuse identity and event inferences, I will focus on three key nodes: opportunity, disposal and guilt.

8.1. Opportunity

To fix some notation, let $\boldsymbol{E_{ID}}$ and $\boldsymbol{E_{SUV}}$ be the identification evidence that supports the proposition ***ID*** that the defendant is the person depicted in the video. Let $\boldsymbol{E_{video}}$ be the evidence that supports the proposition ***N(v)*** that the person depicted in the video was near the crime scene at the relevant time. This way the opportunity hypothesis has been disambiguated into two distinct hypotheses, one about the events (***N(v)***) and another about identity (***ID***). But the hypothesis we care about is ***N(d)***: that the defendant was near the crime scene. How should this be modeled?

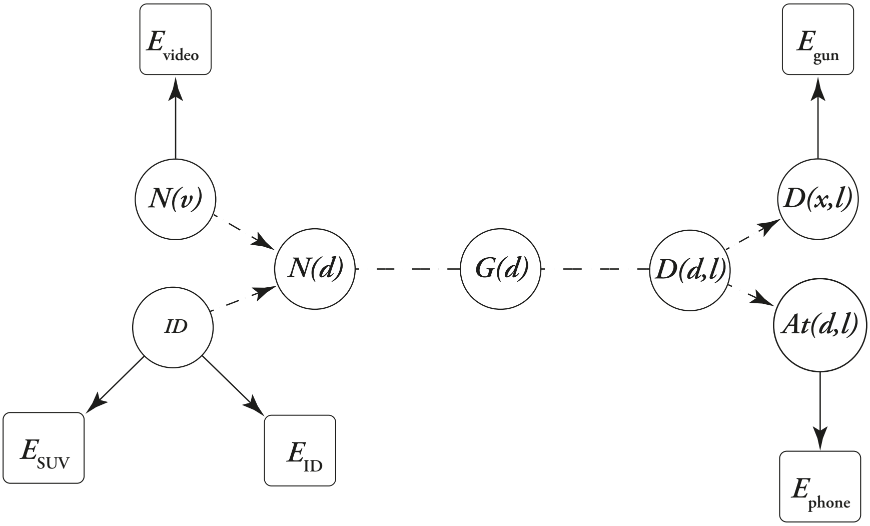

Figure 10. (Left) ***ID*** *is the proposition that the defendant is the person depicted in the video;* ***N(v)*** *is the proposition that the person depicted in the video was near the crime scene. The opportunity proposition* ***N(d)*** *is modeled as depending on* ***ID*** *and* ***Nv****. (Right)* ***D(x, l)*** *is the proposition that a person* ***x*** *disposed of the firearm at location* ***l****;* ***At(d, l)*** *is the proposition that the defendant* ***d*** *was at location* ***l*** *. The dashed arrows indicate analytical dependencies. The dashed lines indicate unspecified dependencies.*

The simplest approach is to think of it as a conjunction:

$$N(d) \coloneqq ID \wedge N(v)$$

But this reading cannot be right. Although the conjunction of the two propositions on the right entails that the defendant was near the crime scene, the converse does not follow. If the defendant was near the crime scene, it does not necessarily follow that he was the person in the video. The defendant might have been near the crime scene in some other way, independently of the events depicted in the surveillance footage. In addition, if ***N(d)*** were a mere conjunction of the other two hypotheses, each item of evidence would behave as though it could make an independent contribution toward supporting ***N(d)***[8]. But, on its own, neither item is probative of

[8] Assume that E_{ID} supports ***ID***, so that
$P(ID \mid E_{ID}) > P(ID)$
Does E_{ID} also support ***N(d)***? Since
$P(N(d) \mid E_{ID}) = P(ID \wedge N(v) \mid E_{ID}) = P(ID \mid E_{ID})\, P(N(v) \mid ID)$
By definition
$P(N(d)) = P(ID)\, P(N(v) \mid ID)$
The question is whether
$P(ID \mid E_{ID})\, P(N(v) \mid ID) > P(ID)\, P(N(v) \mid ID)$
The inequality holds assuming $P(ID \mid E_{ID}) > P(ID)$ unless $P(N(v) \mid ID) = 0$.

the defendant's presence near the scene. As noted before, their evidential force arises only when they are considered together[9].

Here is a better interpretation of the opportunity hypothesis and the inference that supports it. The video evidence supports the proposition that the person shown in the video was near the crime scene (***N(v)***), while the identification evidence supports the proposition that the person shown in the video is the defendant. To make fully transparent the identification between the person ***v*** in the video and the defendant ***d***, let us denote this proposition by ***d = v*** instead of the more generic ***ID***. The inference in question that leads to ***N(d)*** looks like this:

$$\frac{\boldsymbol{N(v)} \qquad \boldsymbol{d=v}}{\boldsymbol{N(d)}}$$

The inference is deductive: once ***N(v)*** and ***d = v*** are accepted, ***N(d)*** follows by substituting co-referring expressions. How does this pattern of inference fit into a standard Bayesian network?

One possible response is to hard code the inference into the conditional probability table for node ***N(d)*** representing the hypothesis that the defendant was near the crime scene. The node will have two parents: ***N(v)*** and ***v=d*** or ***ID*** (Figure 10, left side). The probability table could then be specified as follows:

N(v)	*v = d*	*P(N(d) = 1| N(v), v = d)*
1	1	1
0	1	*q*
1	0	*p*
0	0	*p*

The first row is key. It captures the deductive inference we are after: if the person in the video was near the crime scene, and if that person was the defendant, then the defendant was near the crime scene. In that case, ***N(d)*** follows with probability **1**. The remaining rows assign baseline probabilities ***p*** or ***q***[10].

[9] This problem can perhaps be deflected by noting that the evidential support will often be limited. Suppose ***P*(*ID*) = 1/1 000 000 = $\mathbf{10^{-6}}$** since 1 million people (or whatever the number) could have been in the video, and similarly ***P*(*N(v)*) = $\mathbf{10^{-6}}$** since the video could depict events in many possible locations. Assuming probabilistic independence, ***P*(*ID* ∧ *N(v)*) = $\mathbf{10^{-12}}$**. If the identification evidence E_{ID} makes ***ID*** almost certain, say ***P*(*ID* | E_{ID}) = .99**, then

$$P(N(v) \wedge ID \mid E_{ID}) \approx P(ID \mid E_{ID})\, P(N(v)) = .99 \times 10^{-6} \approx 10^{-6}$$

The probability increase is significant in relative term, but still tiny in absolute. A similar point applies for the video evidence and proposition **(*N(v)*)**.

[10] The second and third rows reflect the idea that, if the person in the video was not the defendant, then facts about the person in the video do not by themselves support or oppose the proposition that the defendant was near the crime scene.

This construction shows that the substitution inference can be represented in a Bayesian network by the appropriate assignment of values in the conditional probability table. The limitation, however, is that the model does not fully express the pattern of inference involved. The underlying reasoning relies on coreference: from $N(x)$ and $x = d$, infer $N(d)$, for any x. A standard (propositional) Bayesian network can encode an instance of the pattern of inference, not the general pattern.

8.2. Disposal

Another modeling challenge arises for the disposal hypothesis. To start, the firearm evidence E_{gun} supports the proposition that someone disposed of a firearm at a definite location. Ballistic evidence—which we have not separately modeled—supports the claim that the firearm recovered at that location was the firearm used in the shooting. So, the firearm evidence support the claim—call it $D(x,l)$ —that someone disposed of the firearm at that specific location. Who did it? The phone location evidence E_{phone} supports the proposition that the defendant was at that same location, call it $At(d,l)$. Given that the location is the same, these propositions—taken together, not independently—support the conclusion $D(d,l)$ that it was the defendant who disposed of the firearm at that location.

As with the inference about opportunity, the conclusion $D(d,l)$ that the defendant disposed of the firearm at that location cannot be modeled as a conjunction of the other two propositions $D(x,l)$ and $At(d,l)$. The inference about the disposal hypothesis can be formalized, as follows:

$$\frac{\exists x D(x, l) \quad At(d, l) \quad \neg Defeater}{D(d, l)}$$

The inference is defeasible because the defendant's presence at l does not deductively establish that he was the person who disposed of the firearm there, although the identity of the locations makes it somewhat likely.

How does this pattern of inference fit into a standard Bayesian network? The proposition $D(d, l)$ that the defendant disposed of the firearm at location l is an explanatory hypothesis. If true, it would make sense of the truth of two otherwise separate intermediate hypotheses: first, that the firearm was disposed of at location l, and second, that the defendant was at that location. So the most natural option is to assign a node for hypothesis $D(d, l)$ with two outgoing arrows pointing to two nodes, one for $\exists x D(x, l)$ and another for $At(d, l)$ (Figure 10, right side). This representation is natural since those hypotheses are two consequences of $D(d, l)$. The following conditional probability tables capture the relationships:

$D(d, l)$	$P(At(d, l) = 1 \mid D(d, l))$
1	1
0	p

$D(d, l)$	$P(d(x, l) = 1 \mid D(d, l))$
1	1
0	q

The first row of each table captures the relevant implication. The second row in each table is a baseline value. If the defendant did not dispose of the weapon at the location, it does not follow that he was not there or that no one disposed of the weapon there. The defendant may have been there for some other reason or someone else may have disposed of the weapon.

This representation does the trick, but again it has a limitation. The probability tables do not explicitly encode the preservation of the same location $\boldsymbol{l}$ across propositions. As with the video and identification inference, the network can hard code the inference, but does not fully capture the general pattern of reasoning.

8.3. Guilt

A further challenge arises from connecting the opportunity and disposal hypotheses—which we have split into identity- and event-level sub-hypotheses—to the guilt hypothesis. It is helpful to distinguish two inferential routes here. First, recall that the video, ID and SUV evidence, together, place the defendant near the crime scene at the relevant time. At the event level, the video evidence supports the further claim that whoever was near the crime scene at that time was very likely the perpetrator. The timing and proximate location are doing much of the work here. The video evidence seems to bear not just on the proposition that someone was near the crime scene, but on the strength of the inference from being near the crime scene to committing the crime. Stated more formally, the first inferential route toward the defendant's guilt is as follows:

$$\frac{\boldsymbol{N(x) \Longrightarrow G(x) \qquad N(d)}}{\boldsymbol{G(d)}}$$

In other words, whoever was near the crime scene at the relevant time was likely the perpetrator, and the defendant was near the crime scene at the relevant time; therefore, the defendant was likely the perpetrator[11].

[11] The arrow$\Longrightarrow$ is meant to signal the defeasible relation between antecedent ***N(x)*** and consequent ***G(x)***. Another approach would use a a conditional probability statement, such as ***P(G(x) | N(x))***, and assign it a reasonably high value.

The second inference toward guilt goes through the disposal route. The phone location evidence links the defendant to the location where the firearm was recovered. The ballistics evidence establishes that the recovered firearm was the firearm used in the shooting. That fact supports a bridge from disposal to guilt: whoever later disposed of that firearm was likely the shooter. Stated more formally,

$$\frac{D(x,l) \Rightarrow G(x) \qquad D(d,l)}{G(d)}$$

The two routes converge on the same conclusion. One moves from the defendant's presence near the crime scene to guilt; the other from the defendant's presence at the location where the firearm was disposed of to guilt.

As before, we could hard code these inferences in a Bayesian network by adding suitable conditional probability tables. But that move would require assigning a defeasible conditional statement to a node in the network. In principle, this is possible: a node can stand for any proposition, however complex. But the cost is that the structure of the inference can no longer be read off the network.

A related oddity concerns the video evidence and the ballistics evidence, which was not formally modeled. These items of evidence appear to support a relation between propositions, rather than simply the truth of a single proposition. The video evidence supports the inference from someone's being near the crime scene at the relevant time to that person's being the shooter. The ballistics evidence supports the inference from someone's disposing of the firearm to that same person's having committed the shooting. In standard propositional Bayesian networks, however, nodes usually represent atomic propositions or events, while inferential relations among propositions are represented through the graph structure and the conditional probability tables. So evidence that bears on the strength of an inferential connection does not have an obvious place in the model [12].

9. CONCLUSION

If the analysis in the previous section is on the right track, the expressive limitation I have identified is not with probability theory as such. The problem lies instead with the expressive resources of standard, propositional Bayesian networks. Their nodes normally stand for atomic propositions, but some legally important inferences have a richer logical structure. The inference about opportunity, for example, depends on identity: if $v = d$, then predicates true of v may be transferred to d. Other

[12] It is in principle possible to add further nodes so that the inferential relation is reified as another proposition. The model can then become formally workable, but such a move would obscure the fact that the evidence is about an inference from one proposition to another.

inferences depend on conditional claims, such as the claim that whoever was near the crime scene at the relevant time was the shooter. As we have seen, a standard, propositional Bayesian network can approximate such inferences by hard coding the relevant relationships into conditional probability tables. But that is only a workaround. Although it is workable in simple cases, one wonders whether this approach breaks down as the legal case to be modeled grows more complex.

A better option would be to rely on a formal framework with greater expressive power. If probability could be combined with a first-order language—made of variables, predicates, conditionals, and identity—then many of the problematic inferences we have examined could be represented more directly. Instead of treating ***v = d***, ***N(v)***, and ***N(d)*** as separate propositional nodes connected by an *ad hoc* table, the model could represent the general identity principle that, if two terms co-refer, predicates true of one are true of the other. Many evidential inferences in legal cases concern people, objects, locations, and events whose identities are uncertain. A satisfactory model should be able to represent inferences as to whether the person in a video is the defendant; whether the firearm recovered at one location is the firearm used in the shooting; whether the person who disposed of the firearm is the same person who appeared near the crime scene; and so on.

The challenge, then, is how probability should be combined with a first-order language. To be sure, there is already a body of work in this direction. Multi Entity Bayesian Networks (MEBNs), for example, were developed as a first-order language for probabilistic reasoning (Laskey, 2008). But MEBNs were not developed for legal applications, so the feasibility of using them to model legal evidence deserves further examination. That is a task for another time.

REFERENCES

Allen, R. J., & Pardo, M. S. (2007). The Problematic Value of Mathematical Models of Evidence. *Journal of Legal Studies*, *36*(1), 107–140.

Bozza, S., Taroni, F., & Aitken, C. (2026). The Myth of Judicial Objectivity: Why Subjectivity is the only Path to a Transparent Trial. *Quaestio Facti. Revista Internacional Sobre Razonamiento Probatorio*, *11*.

Dahlman, C., & Kolflaath, E. (2022). Causal Models versus Reason Models in Bayesian Networks for Legal Evidence. *Synthese*, *200*, 477.

Fenton, N., Neil, M., & Lagnado, D. A. (2013). A General Structure for Legal Arguments about Evidence using Bayesian Networks. *Cognitive Science*, *37*(1), 61–102.

Fenton, N., Neil, M., Yet, B., & Lagnado, D. A. (2020). Analyzing the Simonshaven case using Bayesian Networks. *Topics in Cognitive Science*, *12*(4), 1092–1114.

Hampson, L., & Leeuwen, L. van. (2025). Investigating the Value of Qualitative Bayesian Networks of Complete Cases as "Double-Check" Tools on Traditional Judicial Reasoning: An Exploratory Study. *Proceedings of the Workshop on AI for Evidential Reasoning*, *4157*, 40–71. https://ceur-ws.org/Vol-4157/paper5.pdf

Kadane, J. B., & Schum, D. A. (1996). *A Probabilistic Analysis of the Sacco and Vanzetti Evidence*. John Wiley & Sons.

Lagnado, D. A. (2021). *Explaining the Evidence: How the Mind Investigates the World*. Cambridge University Press.
Laskey, K. B. (2008). MEBN: A Language for First-Order Bayesian Knowledge Bases. *Artificial Intelligence*, *172*(2–3), 140–178.
Mackor, A. R. (2026). Bayesian Modelling of Criminal Cases as a Whole: A Philosophical Reflection on Dutch Case Law. *Quaestio Facti. Revista Internacional Sobre Razonamiento Probatorio*, *10*, 363–384.
Stacey, J., Opperman, S., Curran, J. M., & Harbison, S. A. (2025). Current State and Barriers to Global Adoption of Forensic Evaluative Reporting of Findings Given Activity-Level Propositions. *Science & Justice*, *65*(6), 101336.
Taroni, F., Biedermann, A., Bozza, S., Garbolino, P., & Aitken, C. (2014). *Bayesian Networks for Probabilistic Inference and Decision Analysis in Forensic Science* (2nd ed.). John Wiley & Sons.
Taylor, D., Biedermann, A., Hicks, T., & Champod, C. (2018). A Template for Constructing Bayesian Networks in Forensic Biology Cases when Considering Activity Level Propositions. *Forensic Science International: Genetics*, *33*, 136–146.
Tribe, L. H. (1971). Trial by mathematics: Precision and Ritual in the Legal Process. *Harvard Law Review*, *84*(6), 1329–1393. https://doi.org/10.2307/1339610
Vlek, C. S., Prakken, H., Renooij, S., & Verheij, B. (2014). Building Bayesian Networks for Legal Evidence with Narratives: A Case Study Evaluation. *Artificial Intelligence and Law*, *22*(4), 375–421.