COGNITIVE SCIENCE
A Multidisciplinary Journal

Cognitive Science (2012) 1–42

ISSN: 0364-0213 print / 1551-6709 online
DOI: 10.1111/cogs.12004

A General Structure for Legal Arguments About Evidence Using Bayesian Networks

Norman Fenton,[a] Martin Neil,[a] David A. Lagnado[b]

[a]*School of Electronic Engineering and Computer Science, Queen Mary University of London*
[b]*Cognitive, Perceptual, and Brain Sciences Department, University College London*

Received 22 November 2010; received in revised form 3 May 2012; accepted 4 May 2012

Abstract

A Bayesian network (BN) is a graphical model of uncertainty that is especially well suited to legal arguments. It enables us to visualize and model dependencies between different hypotheses and pieces of evidence and to calculate the revised probability beliefs about all uncertain factors when any piece of new evidence is presented. Although BNs have been widely discussed and recently used in the context of legal arguments, there is no systematic, repeatable method for modeling legal arguments as BNs. Hence, where BNs have been used in the legal context, they are presented as completed pieces of work, with no insights into the reasoning and working that must have gone into their construction. This means the process of building BNs for legal arguments is ad hoc, with little possibility for learning and process improvement. This article directly addresses this problem by describing a method for building useful legal arguments in a consistent and repeatable way. The method complements and extends recent work by Hepler, Dawid, and Leucari (2007) on object-oriented BNs for complex legal arguments and is based on the recognition that such arguments can be built up from a small number of basic causal structures (referred to as *idioms*). We present a number of examples that demonstrate the practicality and usefulness of the method.

Keywords: Legal arguments; Probability; Bayesian networks

1. Introduction

The literature on legal argumentation within legal philosophy and Artificial Intelligence and law is well established (probably dating back to Wigmore, 1913) and extensive—see, for example, Ashley (1990), Bankowski, White, and Hahn (1995), and Prakken (1997). This article is restricted to the role of probabilistic Bayesian reasoning in legal practice, a topic

Correspondence should be addressed to Norman Fenton, School of Electronic Engineering and Computer Science, Queen Mary University of London, London E1 4NS, UK. E-mail: norman@eecs.qmul.ac.uk

that has also been addressed in many articles and books (e.g., Aitken & Taroni, 2004; Dawid, 2002; Evett & Weir, 1998; Faigman & Baglioni, 1988; Fienberg & Schervish, 1986; Finkelstein & Levin, 2001; Friedman, 1987; Good, 2001; Jackson, Jones, Booth, Champod, & Evett, 2006; Matthews, 1997; Redmayne, 1995; Robertson & Vignaux, 1995, 1997; Schum, 2001). What we are especially interested in is the role of such reasoning to improve understanding of legal arguments. For the purposes of this paper, an *argument* refers to any reasoned discussion presented as part of, or as commentary about, a legal case. It is our contention that a Bayesian network (BN), which is a graphical model of uncertainty, is especially well suited to legal arguments. A BN enables us to visualize the causal relationships between different hypotheses and pieces of evidence in a complex legal argument. But, in addition to its powerful visual appeal, it has an underlying calculus (based on Bayes' theorem) that determines the revised probability beliefs about all uncertain variables when any piece of new evidence is presented.

The idea of using BNs for legal arguments is by no means new. Although he referred to the method as ''route analysis,'' what Friedman (1987) proposed was essentially a Bayesian causal graphical approach for reasoning probabilistically about the impact of evidence. Many others (e.g., see Aitken et al., 1995; Dawid & Evett, 1997; Huygen, 2002; Jowett, 2001; Kadane & Schum, 1996; Taroni, Aitken, Garbolino, & Biedermann, 2006; Zukerman & George, 2005; Zukerman, 2010) have explicitly used BNs to model legal arguments probabilistically. Indeed, Edwards (1991) provided an outstanding argument for the use of BNs in which he said of this technology: ''I assert that we now have a technology that is ready for use, not just by the scholars of evidence, but by trial lawyers.'' He predicted such use would become routine within ''2–3 years.'' Unfortunately, he was grossly optimistic for reasons that are fully explained in Fenton and Neil (2011). One of the reasons for the lack of takeup of BNs within the legal profession was a basic lack of understanding of probability and simple mathematics, but Fenton and Neil described an approach (that has recently been used successfully in real trials) to overcome this barrier by enabling BNs to be used without lawyers and jurors having to understand any probability or mathematics. However, while this progress enables nonmathematicians to be more accepting of the results of BN analysis, there is no systematic, repeatable method for modeling legal arguments as BNs. In the many papers and books where such BNs have been proposed, they are usually presented as completed pieces of work, with no insights into the reasoning and working that must have gone into determining why the particular set of nodes and links between them were chosen rather than others. Also, there is very little consistency in style or language between different BN models even when they represent similar arguments. This all means that the process of building a BN for a legal argument is ad hoc, with little possibility for learning and process improvement.

The purpose of this article is, therefore, to show that it is possible to meet the requirement for a structured method of building BNs to model legal arguments. The method we propose complements and extends recent work by Hepler et al. (2007). Their key contribution was to introduce the use of object-oriented BNs as a means of organizing and structuring complex legal arguments. Hepler et al. also introduced a small number of ''recurrent patterns of evidence,'' and it is this idea that we extend significantly in this paper, while accepting the object-oriented structuring as given. We refer to commonly recurring patterns as *idioms*. A set of generic BN idioms was first introduced in Neil, Fenton, and Nielsen (2000). These idioms represented an abstract set of classes of reasoning from which specific cases (called instances) for the problem at hand could be constructed. The approach was inspired by ideas from systems engineering and systems theory and Judea Pearl's recognition that ''Fragmented structures of causal organizations are constantly being assembled on the fly, as needed, from a stock of building blocks'' (Pearl, 1988).

In this article, we focus on a set of instances of these generic idioms that are specific to legal arguments. We believe that the proposed idioms are sufficient in the sense that they provide the basis for most complex legal arguments to be built. Moreover, we believe that the development of a small set of reusable idioms reflects how the human mind deals with complex evidence and inference in the light of memory and processing constraints. The proposed idioms conform to known limits on working memory (Cowan, 2001; Halford, Cowan, & Andrews, 2007; Miller, 1956), and the reusable nature of these structures marks a considerable saving on storage and processing. The hierarchical structuring inherent in the general BN framework also fits well with current models of memory organization (Ericsson & Kintsch, 1995; Gobet et al., 2001; Steyvers & Griffiths, 2008). There is further support for this approach in studies of expert performance in chess, physics, and medical diagnosis, where causal schema and scripts play a critical role in the transition from novice to expert (Chase & Simon, 1973; Ericsson, Charness, Feltovich, & Hoffman, 2006). This fit with the human cognitive system makes the idiom-based approach particularly suitable for practical use by nonspecialists.

In contrast to the object-oriented approach proposed by Hepler et al. (2007), we emphasize the causal underpinnings of the basic idioms. The construction of the BNs always respects the direction of causality, even where the key inferences move from effect to cause. Again this feature meshes well with what is currently known about how people organize their knowledge and draw inferences (Krynski & Tenenbaum, 2007; Lagnado, Waldmann, Hagmayer, & Sloman, 2007; Sloman, 2005; Sloman & Lagnado, 2005). Indeed, the predominant psychological model of legal reasoning, the story model, takes causal schema as the fundamental building blocks for reasoning about evidence (Pennington & Hastie, 1986, 1992). The building block approach means that we can use idioms to construct models incrementally while preserving interfaces between the model parts that ensure they can be coupled together to form a cohesive whole. Likewise, the fact that idioms contain causal information in the form of causal structure alone means any detailed consideration of the underlying probabilities can be postponed until they are needed, or we can experiment with hypothetical probabilities to determine the impact of the idiom on the case as a whole. Thus, the idioms provide a number of necessary abstraction steps that match human cognition and also ease the cognitive burden involved in engineering complex knowledge-based systems.

Bayesian approaches to reasoning and argument are gaining ground in cognitive science (Oaksford & Chater, 2007, 2010). Most relevant to our proposed framework is research by Hahn, Oaksford, and colleagues (Corner, Hahn, & Oaksford, 2011; Hahn, Harris, & Corner, 2009; Hahn & Oaksford, 2007) that proposes a Bayesian account of informal argumentation and argument strength. In particular, Hahn and Oaksford (2007) give a Bayesian analysis of several classical informal reasoning fallacies, including the argument from ignorance, circular, and slippery slope arguments. Although BNs are not a dominant part of their work, a simple network is used to analyze the argument from circularity. This article advances a framework that is consistent with and complementary to this research. It shares the core belief that informal arguments are best analyzed within a Bayesian framework. In contrast to Hahn and Oaksford, our focus is on legal arguments, and BNs play a central role in the proposed framework. We also introduce causal idioms that are tailored to the legal domain and serve as critical building blocks for large-scale legal arguments.

This article is structured as follows: In Section 2, we state our assumptions and notation, while also providing a justification for the basic Bayesian approach. The structured BN idioms are presented in Section 3, while examples of applying the method to complete legal arguments are presented in Section 4. Our conclusions include a roadmap for empirical research on the impact of the idioms for improved legal reasoning. Executable versions of all of the BN models described in the paper are freely available for inspection and use at http://www.eecs.qmul.ac.uk/~norman/Models/legal_models.html.

2. The case for Bayesian reasoning about evidence

We start by introducing some terminology and assumptions that we will use throughout:

1. A legal argument involves a collection of *hypotheses* and *evidence* about these hypotheses.
2. A *hypothesis* is a Boolean statement whose truth value is generally unknowable to a jury. The most obvious example of a hypothesis is the statement ''Defendant is guilty'' (of the crime charged). Any hypothesis like this, which asserts guilt/innocence of the defendant, is called the *ultimate hypothesis*. There will generally be additional types of hypotheses considered in a legal argument, such as ''defendant was present at the crime scene'' or ''the defendant had a grudge against the victim.''
3. A piece of *evidence* is a Boolean statement that, if true, lends support to one or more hypothesis. For example, ''an eye witness testifies that defendant was at scene of crime'' is evidence to support the prosecution hypothesis that ''defendant is guilty,'' while ''an eye witness testifies that the defendant was in a different location at the time of the crime'' is evidence to support the defense hypothesis.
4. We shall assume there is *only one ultimate hypothesis*. This simplifying assumption means that the prosecution's job is to convince the jury that the ultimate hypothesis is true, while the defense's job is to convince the jury it is false. Having a single ultimate hypothesis means that we can use a single argument structure to represent both the prosecution and defense argument.

The situation that we are ruling out for practical reasons is where the defense has at least one hypothesis that is not simply the negation of the prosecution's hypothesis. For example, whereas in a murder case the ultimate hypothesis for the prosecution might be that the defendant is guilty of murder, the defense might consider one or more of the following ultimate hypotheses, none of which is the exact negation of the prosecution's:

1. Defendant is guilty of killing but only in self-defense.
2. Defendant is guilty of killing but due to diminished responsibility.
3. Defendant is guilty of killing but only through hiring a third party who could not be stopped after the defendant changed his or her mind.

If there are genuinely more than one ultimate hypothesis, then a different argument structure is needed for each.

Our approach assumes that the inevitable uncertainty in legal arguments is quantified using probability. However, it is worth noting that some people (including even senior legal experts) are seduced by the notion that ''there is no such thing as probability'' for a hypothesis like ''Defendant is guilty.'' As an eminent lawyer told us: ''Look, the guy either did it or he didn't do it. If he did then he is 100% guilty and if he didn't then he is 0% guilty; so giving the chances of guilt as a probability somewhere in between makes no sense and has no place in the law.'' This kind of argument is based on a misunderstanding of the meaning of uncertainty. Before tossing a fair coin there is uncertainty about whether a ''head'' will be tossed. The lawyer would accept a probability of 50% in this case. If the coin is tossed without the lawyer seeing the outcome, then the lawyer's uncertainty about the outcome is the same as it was before the toss, because he or she has *incomplete information* about an outcome that has happened. The person who tossed the coin knows for certain whether or not it was a ''head,'' but without access to this person the lawyer's uncertainty about the outcome remains unchanged. Hence, probabilities are inevitable when our information about a statement is incomplete. This example also confirms the inevitability of personal probabilities about the same event, which differ depending on the amount of information available to each person. In most cases, the only person who knows for certain whether the defendant is guilty is the defendant. The lawyers, jurors, and judge in any particular case will only ever have partial (i.e., incomplete) information about the defendant's guilt/innocence.

Another common objection to the use of probability theory in legal reasoning, voiced by various legal scholars (see Tillers & Green, 1988), is: Where do the numbers come from? This is an important question, especially when we move from the well-defined examples of dice or coins to the messy real world of crimes and criminals. However, this line of objection often conflates the difficulty of providing precise probabilities with the applicability of the probabilistic framework (Tillers, 2011). The main contribution of probability theory to evidence evaluation is that it provides consistent rules for updating one's beliefs (probabilities) given new evidence. The question of where these initial beliefs come from is a separate issue. Thus, probability theory and BNs, in particular, are predominantly about the structure of probabilistic reasoning, and often the exact probabilities used to analyze a case are not important (and a range of values can be tried out). Moreover, by using the likelihood ratio (see below), which involves the relative comparison between two probabilities, we can evaluate the value of evidence in support of (or against) a hypothesis without having to consider the prior probability of the hypothesis.

Probabilistic reasoning of legal evidence often boils down to the simple causal scenario shown in Fig. 1a (which is a very simple BN): We start with some hypothesis *H* (normally the ultimate hypothesis that the defendant is or is not guilty) and observe some evidence *E* (such as an expert witness testimony that the defendant's blood does or does not match that found at the scene of the crime).

The direction of the causal structure makes sense here because the defendant's guilt (innocence) increases (decreases) the probability of finding incriminating evidence. Conversely, such evidence cannot ''cause'' guilt. Although lawyers and jurors do not formally use Bayes' theorem (and the ramifications of this, for example, in the continued proliferation of probabilistic reasoning fallacies are explained in depth in Fenton & Neil, 2011), they

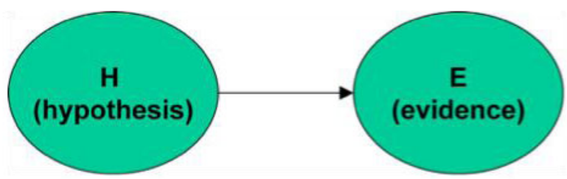

(a)

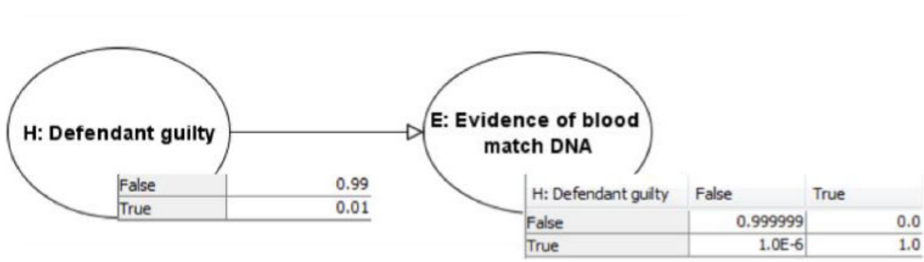

(b)

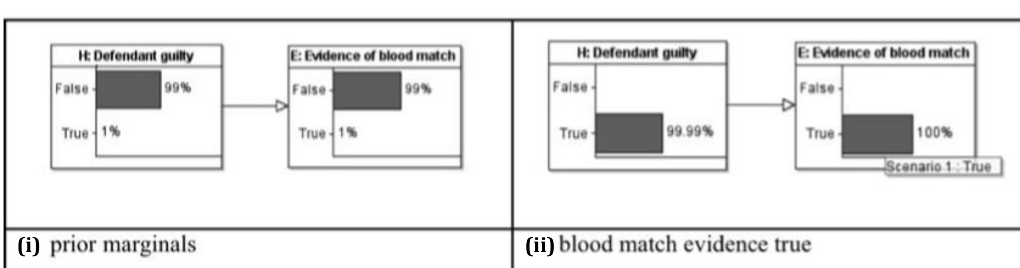

(c)

Fig. 1. (a) Causal view of evidence. (b) Bayesian network (BN) for blood match DNA evidence with Node Probability Tables shown. (c) Running the simple model. Note that in this and all subsequent screenshots of the BN outputs, all probabilities are expressed as percentages rather than values between 0 and 1. Hence, the marginal probability for the defendant being guilty here is *P*(Guilty) = .01 and *P*(not Guilty) = .99.

would normally use the following widely accepted intuitive legal procedure for reasoning about evidence:

1. We start with some (unconditional) prior assumption about guilt (e.g., the ''innocent until proven guilty'' assumption equates to the defendant no more likely to be guilty than any other member of the population).
2. We update our prior belief about H once we observe evidence E. This updating takes account of the *likelihood* of the evidence, which is the chance of seeing the evidence E if H is true.

This turns out to be a perfect match for Bayesian inference. Formally, we start with a prior probability $P(H)$ for the hypothesis H; the likelihood, for which we also have prior knowledge, is formally the conditional probability of E given H, which we write as $P(E|H)$. Bayes' theorem provides the formula for updating our prior belief about H in the light of observing E. In other words, Bayes calculates $P(H|E)$ in terms of $P(H)$ and $P(E|H)$. Specifically,

$$P(H|E) = \frac{P(E|H)P(H)}{P(E)} = \frac{P(E|H)P(H)}{P(E|H)P(H) + P(E|\text{not}H)P(\text{not}H)}$$

As an example, assume for simplicity that a blood trace found at the scene of a crime must have come from the person who committed the crime. The blood is tested against the DNA of the defendant and the result (whether true or false) is presented. This is certainly an important piece of evidence that will adjust our prior belief about the defendant's guilt. Using the approach described above, we could model this using the BN shown in Fig. 1b where the tables displayed are the node probability tables (NPTs) that are specified as prior beliefs.

Here, we have assumed that the ''random DNA match probability'' is one in a million, which explains the entry in the NPT for $P(E \mid \text{not } H)$ (the probability of a match in an innocent person). We also assume, for simplicity, that we will definitely establish a match if the defendant is guilty, that is, $P(E|H) = 1$, and that the DNA analysis procedures are perfect (see Fenton & Neil, 2012 for a full discussion of the implications when these assumptions do not hold). Finally, we have assumed that the prior probability of guilt is .01 (which would be the case if, e.g., the defendant was one of 100 people at the scene of the crime). With these assumptions, the marginal distributions (i.e., the probabilities before any evidence is known) are shown in Fig. 1c (i).

If we discover a match then, as shown in Fig. 1c (ii), when we enter this evidence, the revised probability for guilt jumps to 99.99%, that is, .9999, so the probability of innocence is now one in 10,000. Note that although this is a small probability, it is still significantly greater than the random match probability; confusing these two is a classic example of the prosecutor's fallacy (Fenton & Neil, 2011). There are, of course, a number of simplifying assumptions in the model here that we will return to later. To avoid fundamental confusions, a number of key points about this approach need to be clarified.

2.1. The inevitability of subjective probabilities

Ultimately, any use of probability—even if it is based on frequentist statistics—relies on a range of subjective assumptions. Hence, it is irrational to reject the principle of using subjective probabilities. The objection to using subjective priors may be calmed in many cases by the fact that it may be sufficient to consider a range of probabilities, rather than a single value for a prior. For example, in the real case described in Fenton and Neil (2010) it was shown that taking both the most pessimistic and most optimistic priors, when the impact of the evidence was considered, the range of the posterior probabilities always comfortably pointed to a conclusive result for the main hypothesis.

2.2. Avoiding dependence on prior probabilities by using the ''likelihood ratio''

It is possible to avoid the delicate and controversial issue of assigning a subjective prior probability to the ultimate hypothesis (or indeed to any specific hypothesis) if we instead are prepared to focus on the probabilistic ''value'' of the evidence. Specifically, the value of any single piece of evidence E on a hypothesis H can be determined by considering only the *likelihood ratio* of E. Informally, the likelihood ratio for E tells us how much more likely we are to see the evidence E if the prosecution hypothesis is true compared to if the defense hypothesis is true. Formally, it is the probability of seeing the evidence E if H is true (e.g., ''defendant is guilty'') divided by the probability of seeing that evidence if H is not true (e.g., ''defendant is not guilty''), that is, $P(E|H)$ divided by $P(E \mid \text{not } H)$. For example, in the case of the DNA evidence above, the likelihood ratio is one million, since $P(E|H) = 1$ and $P(E \mid \text{not } H) = .0000001$.

An equivalent form of Bayes' theorem (called the ''odds'' version of Bayes') provides us with a concrete meaning for the likelihood ratio. Specifically, this version of Bayes tells us that the posterior odds of H are the prior odds times the likelihood ratio (see Pearl, 1988 and Fenton, 2011 for further details). So, if the likelihood ratio is 1 million (as in our DNA example), this means that whatever the prior odds were in favor of guilt, the posterior odds must increase by a factor of 1 million as a result of seeing the evidence. So, if our prior belief was that the odds were a million to one *against* guilty, then after the seeing the evidence the odds swing to ''evens''; but if our prior belief was that the odds were a only 10 to 1 against guilty, then after the seeing the evidence the odds swing to 100,000 to 1 in favor of guilty. In general, if the likelihood ratio is bigger than 1, then the evidence increases the probability of H (with higher values leading to higher probability of guilt), while if it is <1, it decreases the probability of H (and the closer it gets to zero the lower the probability of H). If the likelihood ratio is equal (or close) to 1, then E offers no real value at all as it neither increases nor decreases the probability of guilt. Thus, for example, Evett's crucial expert testimony in the appeal case of Barry George (R v George, 2007), previously convicted of the murder of the TV presenter Jill Dando, focused on the fact that the forensic gunpowder evidence that had led to the original conviction actually had a likelihood ratio of about 1. This is because both $P(E \mid \text{Guilty})$ and $P(E \mid \text{not Guilty})$ were approximately equal to .01. Yet, only $P(E \mid \text{not Guilty})$ had been presented at the original trial.

Although the likelihood ratio enables us to assess the impact of evidence on H without having to consider the prior probability of H, it is clear from the above DNA example that the prior probability must ultimately be considered before returning a verdict, since even knowing that the odds in favor of guilt increase by a factor of 1 million may not ''prove guilt beyond reasonable doubt'' if this is the only evidence against the defendant. That is because we already assume intuitively in such circumstances that the prior probability of guilt is also very low. However, with or without a Bayesian approach, jurors inevitably have to make these considerations. A key benefit of the Bayesian approach is to make explicit the ramifications of different prior assumptions. So a judge could state something like: ''Whatever you believed before about the possible guilt of the defendant, the evidence is one million times more likely if the defendant is guilty than if he is innocent. So, if you believed at the beginning that there was a 50:50 chance that the defendant was innocent, then it is only rational for you to conclude with the evidence that there is only a million to one chance the defendant really is innocent. On this basis you should return a guilty verdict. But if you believed at the beginning that there are a million other people in the area who are just as likely to be guilty of this crime, then it is only rational for you to conclude from the evidence that there is a 50:50 chance the defendant really is innocent. On that basis you should return a not guilty verdict.'' Note that such an approach does not attempt to force particular prior probabilities on the jury (the judiciary would always reject such an attempt)—it simply ensures that the correct conclusions are drawn from what may be very different subjective priors.

Although the examples in the rest of this article *do* consider the prior probability for a hypothesis H and compare this with the posterior probability once the evidence is observed, we could equally as well have produced the likelihood ratio for the evidence. To do this, we would choose any prior (such as assigning equal probability to H being true and false) and then divide the posterior odds for H by the chosen prior odds for H.

2.3. The importance of determining the conditional probabilities in an NPT

When lay people are first introduced to BNs, there is a tendency to recoil in horror at the thought of having to understand and/or complete an NPT such as the one for $E|H$ in Fig. 1b. But, in practice, the very same assumptions that are required for such an NPT are normally made implicitly anyway. The benefit of the NPT is to make the assumptions explicit rather than hidden.

2.4. The need to leave the Bayesian calculations to a Bayesian calculator

Whereas Fig. 1 models the simplest legal argument (a single hypothesis and a single piece of evidence), we generally wish to use BNs to model much richer arguments involving multiple pieces of possibly linked evidence. While humans (lawyers, police, jurists, etc.) must be responsible for determining the prior probabilities (and the causal links) for such arguments, it is simply wrong, as argued in Fenton and Neil (2011), to assume that humans must also be responsible for understanding and calculating the revised probabilities that result from observing evidence. For example, even if we add just two additional pieces of evidence to get a BN like the one in Fig. 2, the calculations necessary for correct Bayesian inference become extremely complex. However, while the Bayesian calculations quickly become impossible to do manually, any Bayesian network tool (e.g., AgenaRisk, 2012; Hugin, 2011) enables us to do these calculations instantly.

Despite its elegant simplicity and natural match to intuitive reasoning about evidence, practical legal arguments normally involve multiple pieces of evidence (and other issues) with complex causal dependencies. This is the rationale for the work begun in Hepler et al. (2007) and that we now extend further by showing that there are unifying underlying concepts which mean we can build relevant BN models, no matter how large, that are still conceptually simple because they are based on a very small number of repeated ''idioms'' (where an idiom is a generic BN structure). We present these crucial idioms in the next section.

3. The idiom-based approach

The application of BNs to real-world domains involves various challenges, including the extension to large-scale problems and the provision of principled guidelines for BN construction. To address these issues, Neil et al. (2000) presented five idioms that cover a wide range of modeling tasks (see Fig. 3).

1. *Cause-consequence idiom* (*Fig. 3a and b*)—models the uncertainty of a causal process with observable consequences. Such a process could be physical or cognitive. This idiom is used to model a process in terms of the relationship between its causes (those events or facts that are inputs to the process) and consequences (those events or factors that are outputs of the process). The causal process itself can involve transforming an existing input into a changed version of that input or by taking an

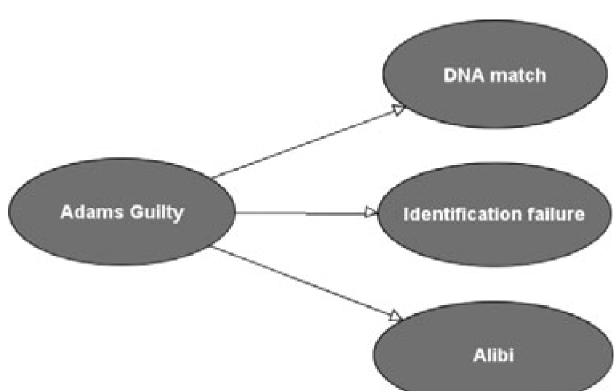

Fig. 2. Hypothesis and evidence in the case of R v Adams (as discussed in Dawid, 2002).

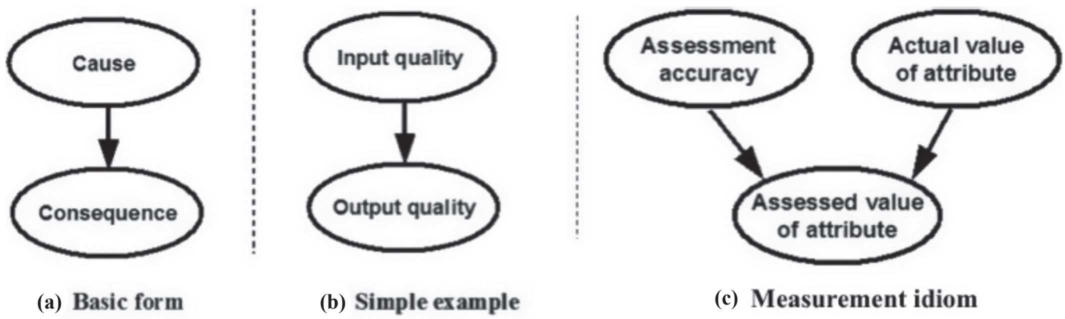

Fig. 3. Generic idioms from Neil et al. (2000).

input to produce a new output. A causal process can be natural, mechanical, or mental in nature. The cause-consequence idiom is organized chronologically—the parent nodes (inputs) can normally be said to come before (or at least contemporaneously with) the children nodes (outputs). Likewise, support for any assertion of causal reasoning relies on the premise that manipulation or change in the causes affects the consequences in some observable way.

2. *Measurement idiom* (*Fig. 3c*)—models the uncertainty about the accuracy of some measurement. We use this idiom to reason about the uncertainty we may have about our own judgments, those of others, or the accuracy of the instruments we use to make measurements. The measurement idiom represents uncertainties we have about the process of observation. By observation, we mean the act of determining the true attribute, state, or characteristic of some entity. The causal directions here can be interpreted in a straightforward way. The true (actual) value must exist before the observation in order for the act of measurement to take place. Next, the measurement instrument interacts (physically, functionally, or cognitively) with the entity under evaluation and produces some result. This result can be more or less accurate depending on intervening circumstances and biases.
3. *Definitional idiom*—models the formulation of many uncertain variables that together form a functional, taxonomic, or an otherwise deterministic relationship.
4. *Induction idiom*—models the uncertainty related to inductive reasoning based on populations of similar or exchangeable members.
5. *Reconciliation idiom*—models the reconciliation of results from competing measurement or prediction systems.

In this article, we are primarily interested in instances of the cause-consequence, measurement, and definitional idioms.

3.1. Idioms for legal reasoning

As noted in the introduction, a major obstacle to the application of BNs to legal arguments is the lack of principled guidelines for model construction. Although BNs have been discussed in the context of legal arguments (several references are provided in the Introduction), there is no systematic method for modeling legal arguments as BNs. Hence, where BNs have been used in the legal context, they are presented as completed pieces of work, with no insights into the reasoning and working that have gone into their construction. This means the process of building BNs for legal arguments is ad hoc, with little possibility for learning and process improvement. To address this problem, we introduce an idiom-based method for building legal arguments in a consistent and repeatable way.

This proposal adapts and extends the idioms introduced by Neil et al. (2000) and is consistent with the related notion of argumentation schemes (Walton, Reed, & Macagno, 2008). However, the latter approach differs from our proposal in several important ways. Argumentation schemes aim to cover a very broad range of reasoning patterns (almost 100 different schemes are proposed) and do not focus on legal arguments in particular. More important, these schemes explicitly avoid the use of probabilities or BNs, and instead adopt simple nonprobabilistic rules for argument evaluation (for details, see Walton, 2008). However, even in simple cases, these rules can yield evaluations that are contrary to the laws of probability and also conflict with people's intuitive evaluations (see Hahn, Oaksford, & Harris, in press).

3.2. The evidence idiom

We can think of the simple BN in Fig. 1 (and its extension to multiple pieces of evidence in Fig. 2) as the most basic BN idiom for legal reasoning. This basic idiom, which we call the evidence idiom, is an instantiation of the cause-consequence idiom and has the generic structure shown in Fig. 4.

We do not distinguish between evidence that supports the prosecution (*H* true) and evidence that supports the defense (*H* false) as the BN model handles both types of evidence seamlessly. Hence, this idiom subsumes two of the basic patterns in Hepler et al. (2007), namely:

1. *Corroboration pattern*: this is simply the case where there are two pieces of evidences E1 and E2 that both support one side of the argument.

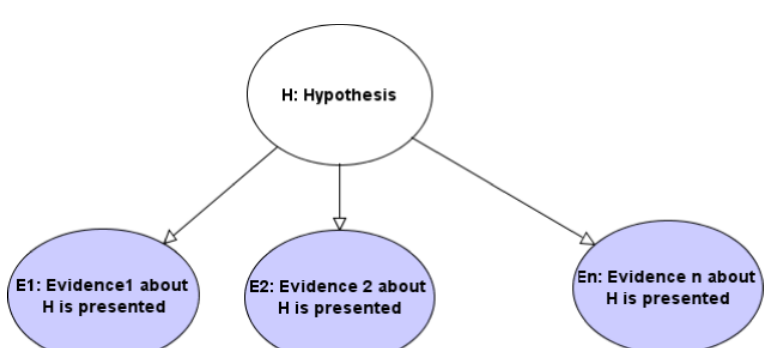

Fig. 4. Evidence idiom.

2. *Conflict pattern*: this is simply the case where there are two pieces of evidences E1 and E2 with one supporting the prosecution and the other supporting the defense.

The evidence idiom has a number of limitations in real cases. The following idioms identify and address these various limitations in turn.

3.3. *The evidence-accuracy idiom*

Let us return to the example of Fig. 1b of evidence in the form of matching DNA from blood found at the scene of the crime. It turns out that the simple model presented made all of the following previously unstated assumptions:

1. The blood tested really was that found at the scene of the crime.
2. The blood did not become contaminated at any time.
3. The DNA testing is perfect; in particular, there is no possibility of wrongly finding a match (note that this is very different to the assumption inherent in the random match probability).
4. The person presenting the DNA evidence in court does so in a completely truthful and accurate way.

If any of the above is uncertain (which may be the case even for DNA evidence, as shown for example in Dror & Hampikian, 2011; Thompson, 2009), then the presentation of evidence of blood match DNA being true or false cannot be simply accepted unconditionally. It must necessarily be conditioned on the overall *accuracy/reliability* of the evidence. In general, the validity of any piece of evidence has uncertainty associated with it, just as there is uncertainty associated with the main hypothesis of guilt. A more appropriate model for this example is therefore the one presented in Fig. 5, which is an instantiation of the measurement idiom.

For simplicity, we have lumped together all possible sources of inaccuracy into a single node (we shall consider a more complete solution later). Because we have introduced a new variable *A* into the model, the NPT for the node *E* is more complex. We can think of the original model as being a special case of this model where *A* was never a doubt (i.e., the accuracy of the evidence was always ''true''). So when *A* is true, the NPT for the node *E* is identical to the NPT in the original model. What is different about the NPT as specified in Fig. 5a is the inclusion of our assumptions about the probability of *E* when *A* is false.

The initial probabilities are shown in Fig. 5b. When evidence of a blood match is presented (Fig. 5c), the probability of guilty increases from the prior 1% to just over 16%. Those who are new to Bayesian reasoning may be surprised that the probability of guilt is so low despite the very low (one in a million) random match probability error.

In fact, the model is working rationally because it is looking for the most likely explanation of the blood match evidence. The prior probability of guilt was 1 in a 100, and this is low compared to the prior probability of inaccurate evidence (1 in 10). So, when only the blood match evidence is presented, the model points to inaccurate evidence as being a more likely explanation for the result. Indeed, the probability of inaccurate evidence jumps from 10% to nearly 85%.

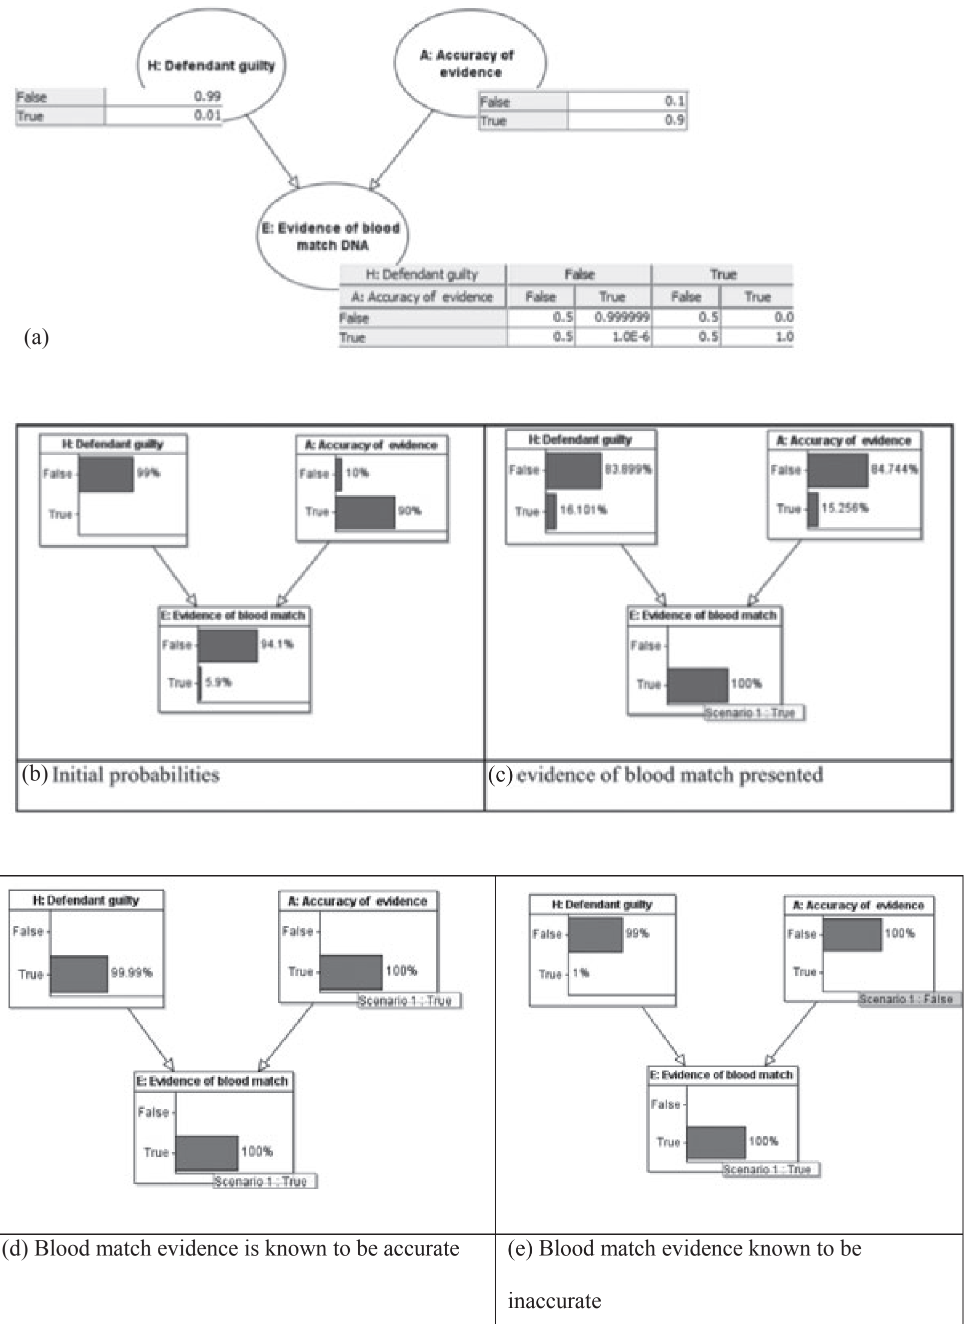

Fig. 5. (a) Revised model: Evidence conditioned on its accuracy with Node Probability Tables shown. (b) Running the model: Initial probabilities. (c) With evidence of blood match. (d) Blood match known to be accurate. (e) Blood match known to be inaccurate.

However, if we determine that the evidence is accurate, as shown in Fig. 5d, the probability of guilt now jumps to 99.99%—the same result as in Fig. 1c because in this scenario the same assumptions are being made. This is an example of ''explaining away'' evidence. If we determine the evidence is inaccurate, the result is shown in Fig. 5e. In this case, the evidence is worthless and the probability of guilt is unchanged from its prior value of 1 in a 100.

By explicitly representing evidence accuracy with a separate variable in the BN, it is much easier to see that the prior probabilities of both guilt and evidence accuracy are relevant to computing the probability of guilt given the evidence report (DNA match). More generally, this idiom clarifies what inferences should be drawn from a positive test result. This is of practical importance because people (including medical experts) are notoriously poor at calculating the true impact of positive test results (Casscells, Schoenberger, & Graboys, 1978; Kahneman, Slovic, & Tversky, 1982). A common error is to ignore the prior probabilities (base-rate neglect) and assume that the probability of the hypothesis (diagnosis) given the evidence is equivalent to the probability of the evidence given the hypothesis (akin to the prosecutor's fallacy; see Balding & Donnelly, 1994). Use of the BN idiom is likely to reduce this error, by making the problem structure explicit. Indeed, a recent set of empirical studies (Krynski & Tenenbaum, 2007) shows that base-rate neglect is attenuated when people have an appropriate causal model on which to map the statistics. This supports the use of causal idioms for rational inference.

The general idiom to model evidence accuracy is shown in Fig. 6. It is an instance of the measurement idiom because we can think of the evidence as simply a measure of (the truth of) the hypothesis. The more accurate the evidence, the closer the evidence value is to the real truth value of the hypothesis. This approach to modeling the accuracy of evidence reports has also been proposed by Bovens and Hartmann (2003). They are primarily concerned with issues in epistemology and the philosophy of science, and they apply BNs to the general context of source reliability. They introduce a reliability node that is essentially equivalent to the accuracy node proposed here. They apply this analysis to a range of problems, including testimonial evidence from multiple sources and scientific hypothesis testing

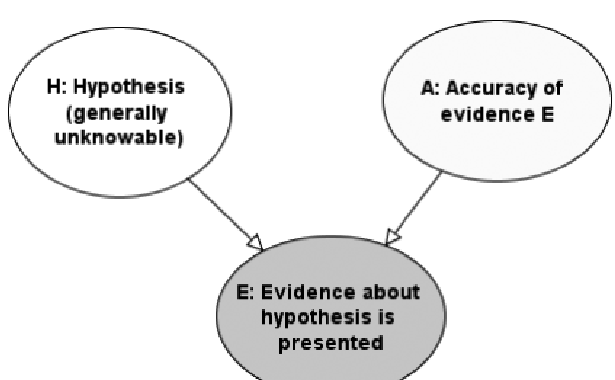

Fig. 6. General idiom to model evidence taking account of its accuracy.

with partially reliable instruments. Their use of BNs to model uncertain information is very consistent with the approach adopted in this article.

To take account of all the individual sources of uncertainty for the DNA blood match example explained at the start of the section, we simply apply the idiom repeatedly as shown in Fig. 7 (of course the different accuracy nodes will in general have different prior probabilities).

There are a number of ways in which the evidence accuracy can be tailored. In particular,

1. There is no need to restrict the node *accuracy of evidence* to being a Boolean (false, true). In general, it may be measured on a more refined scale, for example, a ranked scale (very low, low, medium, high, very high) where very low means ''completely inaccurate'' and very high means ''completely accurate'' or even a continuous scale (although the latter requires special BN algorithms and tools that implement them—- see Neil, Tailor, & Marquez, 2007).

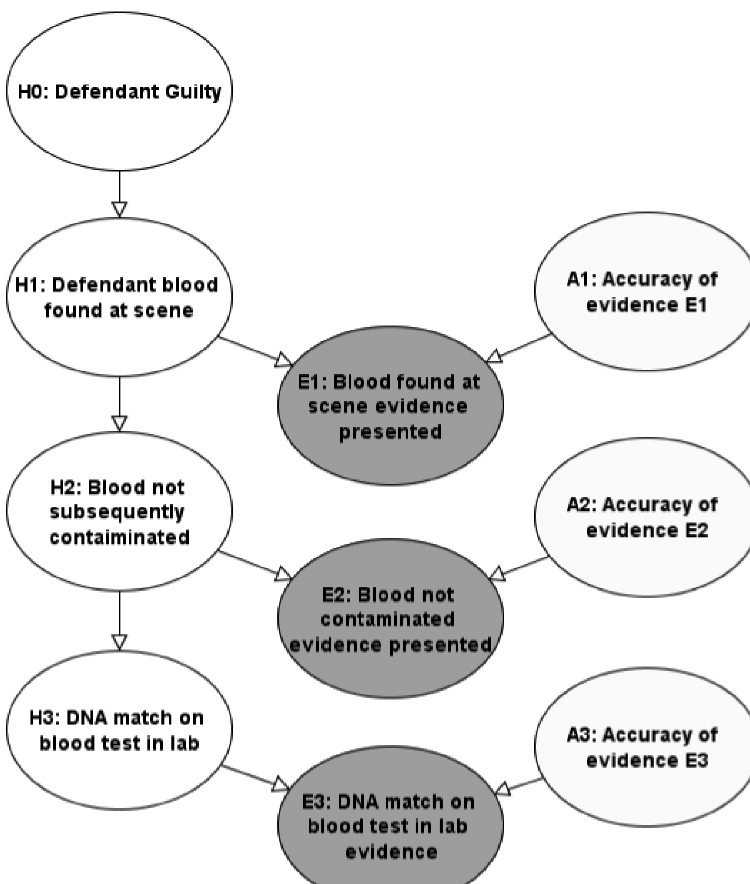

Fig. 7. Full Bayesian network for DNA blood match evidence accuracy.

2. In the case of eyewitness evidence, it is possible to extend the idiom by decomposing ''accuracy'' into three components: *competence*, *objectivity*, and *veracity* as shown in Fig. 8a.

This is essentially what is proposed in Hepler et al. (following on from Schum, 2001), who use the word ''credibility'' to cover what we call ''accuracy,'' although it should be noted that they use an unusual causal structure in which competence influences objectivity, which in turn influences veracity. Our decomposition of accuracy is simply an instance of the definitional idiom. This version of the idiom could also be represented using the object-oriented notation used in Hepler et al. (2007); this is shown in Fig. 8b.

(a)

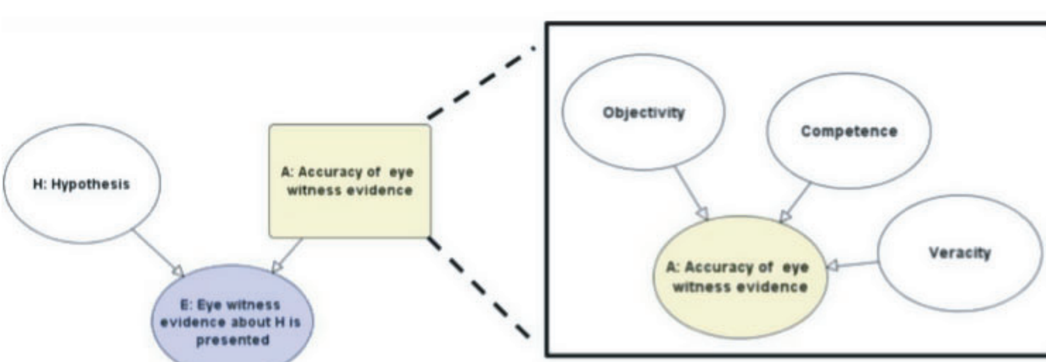

(b)

Fig. 8. (a) Eyewitness evidence-accuracy idiom. (b) Eyewitness accuracy idiom shown using object-oriented structuring.

3.4. Idioms to deal with the key notions of ''motive'' and ''opportunity''

In the examples so far, the ultimate hypothesis (defendant is guilty) has been modeled as a node with no parents. As discussed, this fits naturally with the intuitive approach to legal reasoning whereby it is the hypothesis about which we start with an unconditional prior belief before observing evidence to update that belief. However, there are two very common types of evidence which, unlike all of the examples seen so far, support hypotheses that are *causes*, rather than *consequences*, of guilt. These hypotheses are concerned with ''opportunity'' and ''motive,'' and they inevitably change the fundamental structure of the underlying causal model.

3.4.1. Opportunity

When lawyers refer to ''opportunity'' for a crime, they actually mean a necessary requirement for the defendant's guilt. By far, the most common example of opportunity is ''being present at the scene of the crime.'' So, for example, if Joe Bloggs is the defendant charged with slashing the throat of Fred Smith at four Highlands Gardens on 1 January 2011, then Joe Bloggs had to be present at four Highlands Gardens on 1 January 2011 to be guilty of the crime. The correct causal BN model to represent this situation (incorporating the evidence-accuracy idiom) is shown in Fig. 9a.

Note that, just as the hypothesis ''defendant is guilty'' is unknowable to a jury, the same is true of the opportunity hypothesis. Just like any hypothesis in a trial, its truth value must be determined on the basis of evidence. In this particular example, there might be multiple types of evidence for the opportunity hypothesis, each with different levels of accuracy as shown in Fig. 9b.

From a Bayesian inference perspective, the explicit introduction of opportunity into a legal argument means that it is no longer relevant to consider the prior unconditional probability of the ultimate hypothesis (defendant guilty). Although this destroys the original simplified approach, it does actually make the overall demands on both the jury and lawyers much clearer as follows:

1. The hypothesis requiring an unconditional prior now is that of the opportunity. Unlike the ultimate hypothesis, it is much more likely to be able to base the prior for opportunity on objective information such as the proximity of the defendant's work/-home and the frequency with which the defendant was previously present at the location of the crime scene.
2. Determining the NPT for the conditional probability of the ultimate hypothesis given the opportunity (i.e., *H2|H1*) also forces the lawyers and jurors to consider rational information such as the total number of people who may have been present at the crime scene.

3.4.2. Motive

There is a widespread acceptance within the police and legal community that a crime normally requires a motive (this covers the notions of ''intention'' and ''premeditation'').

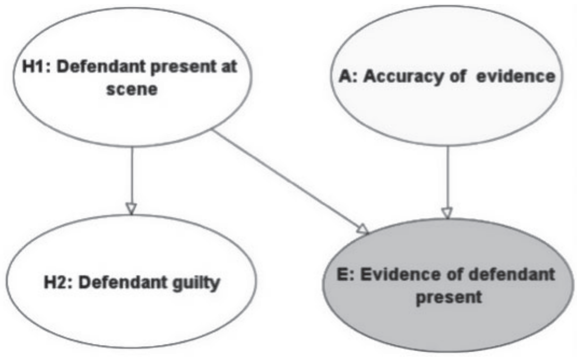

(a)

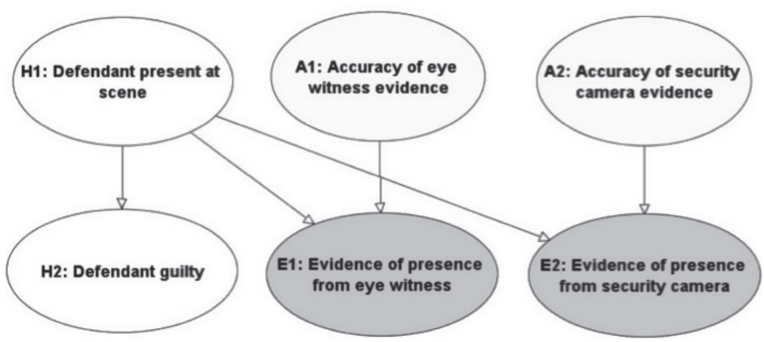

(b)

Fig. 9. (a) Idiom for incorporating ''opportunity'' (defendant present at scene of crime). (b) Multiple types of evidence for opportunity hypothesis.

Although, unlike opportunity, a motive is not a necessary requirement for a crime, the existence of a motive increases the chances of it happening. This means that, as with opportunity, the correct causal BN model to represent motive in a legal argument is shown in Fig. 10.

As with opportunity, the introduction of a motive into a legal argument means that it is no longer relevant to consider the prior unconditional probability of the ultimate hypothesis (defendant guilty). But again, this actually makes the overall demands on both the jury and lawyers much clearer as follows:

1. Although determining an unconditional prior for motive may be just as hard as determining an unconditional prior for guilt, the argument will in general not be so sensitive to the prior chosen. This is because a motive will generally only be introduced if the lawyer has strong evidence to support it, in which case, irrespective of the prior, its truth value will generally be close to true once the evidence is presented.
2. Hence, what really matters is determining the conditional probability of the ultimate hypothesis given the motive. Making this explicit potentially resolves what many

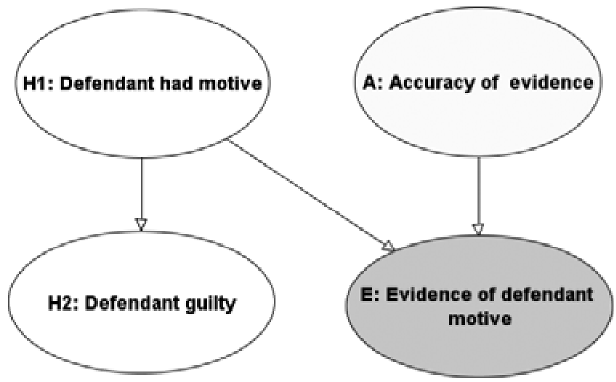

Fig. 10. Idiom for incorporating motive.

believe is one of the most confusing aspects of any trial. Indeed, any lawyer who introduces the notion of a motive should be obliged to state what he or she believes the impact of that motive on guilt to be.

If we wish to include both opportunity and motive into the argument, then the appropriate BN idiom is shown in Fig. 11.

This makes the task of defining the NPT for the ultimate hypothesis *H* a bit harder, as we must consider the probability of guilt conditioned on both opportunity and motive, but again these specific conditional priors are inevitably made implicitly anyway.

What we *do* need to avoid is conditioning *H* directly on *multiple* motives, that is, having multiple motive parents of *H* as shown in Fig. 12a. Instead, if there are multiple motives, we simply model what the lawyers do in practice in such cases: specifically, they consider the accuracy of each motive separately but jointly think in terms of the strength of overall motive. The appropriate model for this is shown in Fig. 12b (using the object-oriented notation). When expanded, with the accuracy nodes included, we get the full model shown in Fig. 13.

It is worth noting that Hepler et al. (2007) introduce nodes for both motive and opportunity. However, they do not consider them as special idioms; instead, they treat both of these the same as any other evidence about the guilty hypothesis, that is, the links are from guilty

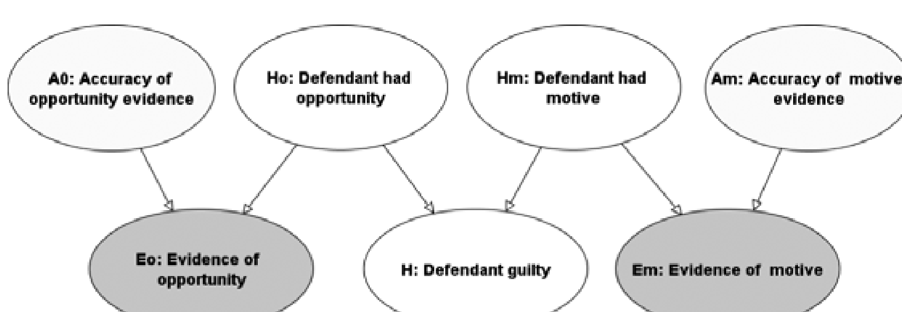

Fig. 11. Bayesian network incorporating both opportunity and motive.

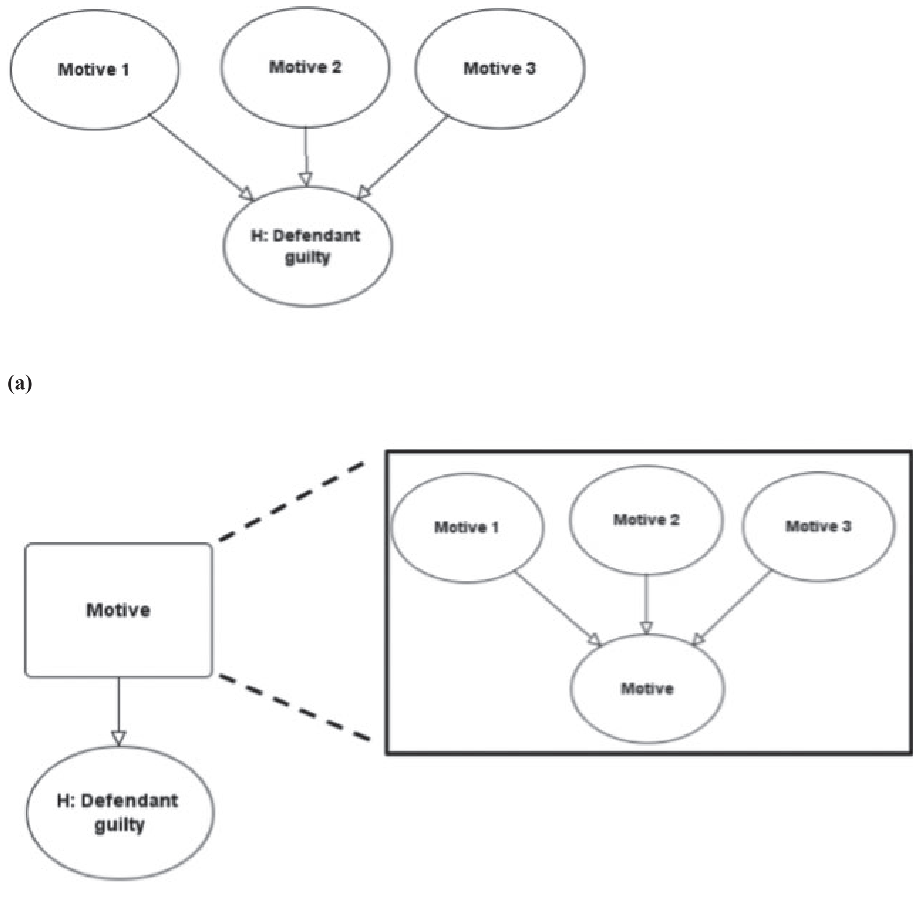

Fig. 12. (a) A structure to be avoided, conditioning *H* on multiple motives. (b) Appropriate model for multiple motives (using object-oriented notation).

to motive and guilty to opportunity rather than the other way round. We believe that this is both structurally wrong and incompatible with standard legal reasoning. Motive and opportunity are typically preconditions for guilt and thus should be modeled as causes (parents) rather than effects of guilt. This proves especially important when more complex combinations of evidence are modeled. For instance, evidence for motive or opportunity occupies a different structural position from direct evidence for guilt (see section below on the difference between direct and circumstantial evidence).

3.5. Idiom for modeling dependency between different pieces of evidence

In the case of a hypothesis with multiple pieces of evidence (such as in Fig. 13), we have so far assumed that the pieces of evidence were *independent* (conditional on *H*). But, in

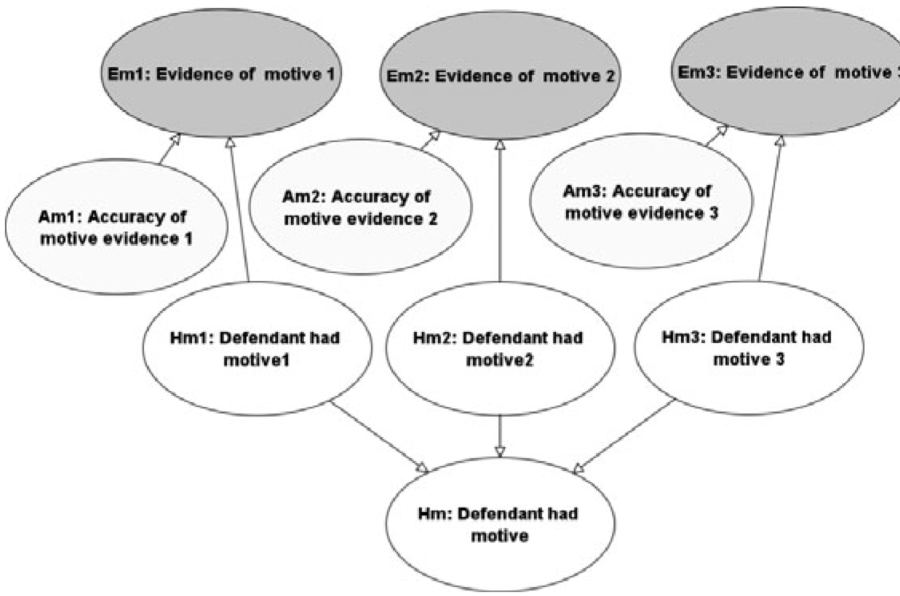

Fig. 13. Full expanded model for multiple motives.

general, we cannot make this assumption. Suppose, for example, that the two pieces of evidence for ''defendant present at scene'' were images from two video cameras. If the cameras were of the same make and were pointing at the same spot, then there is clear dependency between the two pieces of evidence: If we know that one of the cameras captures an image of a person matching the defendant, there is clearly a very high chance that the same will be true of the other camera, irrespective of whether the defendant really was or was not present. Conversely, if one of the cameras does not capture such an image, there is clearly a very high chance that the same will be true of the other camera, irrespective of whether the defendant really was not present. The appropriate way to model this would be as shown in Fig. 14a (for simplicity, we are ignoring the issue of accuracy here) with a direct dependency between the two pieces of evidence. Also, for simplicity, note from the NPTs that ''dependence'' here means the cameras will produce identical results (we can easily adjust the NPT to reflect partial dependence by, for example, making the probability .9 [as opposed to 1] that camera 2 will return ''true'' when *H* is true and camera 1 is true).

If we assume that the prior for *H* being true is 0.1 and the prior for the cameras being dependent is 0.5, then the initial marginal probabilities are shown in Fig 14b. It is instructive to compare the results between the two models: (a) where no direct dependence between E1 and E2; and (b) where it is. Hence, in Fig. 15a and b, we show both these cases where evidence E1 is true. Although both models result in the same (increased) revised belief in *H*, the increased probability that E2 will also be true is different. In (a), the probability increases to 43%, but in (b), the probability is 100% as here we know E2 will replicate the result of E2.

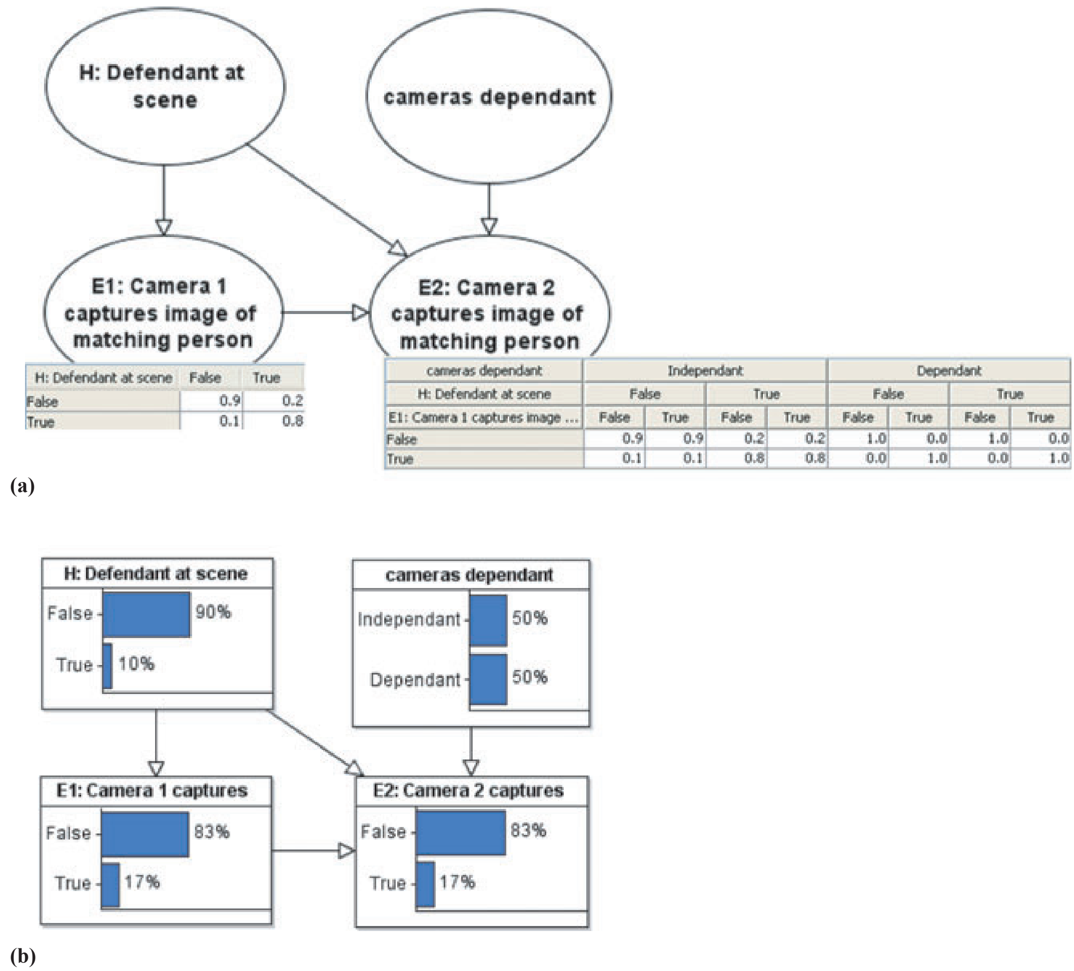

Fig. 14. (a) Idiom for modeling dependency between different pieces of evidence. (b) Running model with initial probabilities.

Fig. 15c and d shows the results of E1 and E2 being presented as true in both cases. When they are dependent, the additional E2 evidence adds no extra value. However, when they are independent, our belief in *H* increases to 88%.

The benefits of making explicit the direct dependence between evidence are enormous. For example, in the case of the Levi Bellfield trial (described in Fenton & Neil, 2011), the prosecution presented various pieces of directly dependent evidence in such a way as to lead the jury to believe that they were independent, hence drastically overstating the impact on the hypothesis being true. In fact, a simple model of the evidence based on the structure above showed that, once the first piece of evidence was presented, the subsequent evidence was almost useless, in the sense that it provided almost no further shift in the hypothesis probability. A similar problem of treating dependent evidence as independent was a key

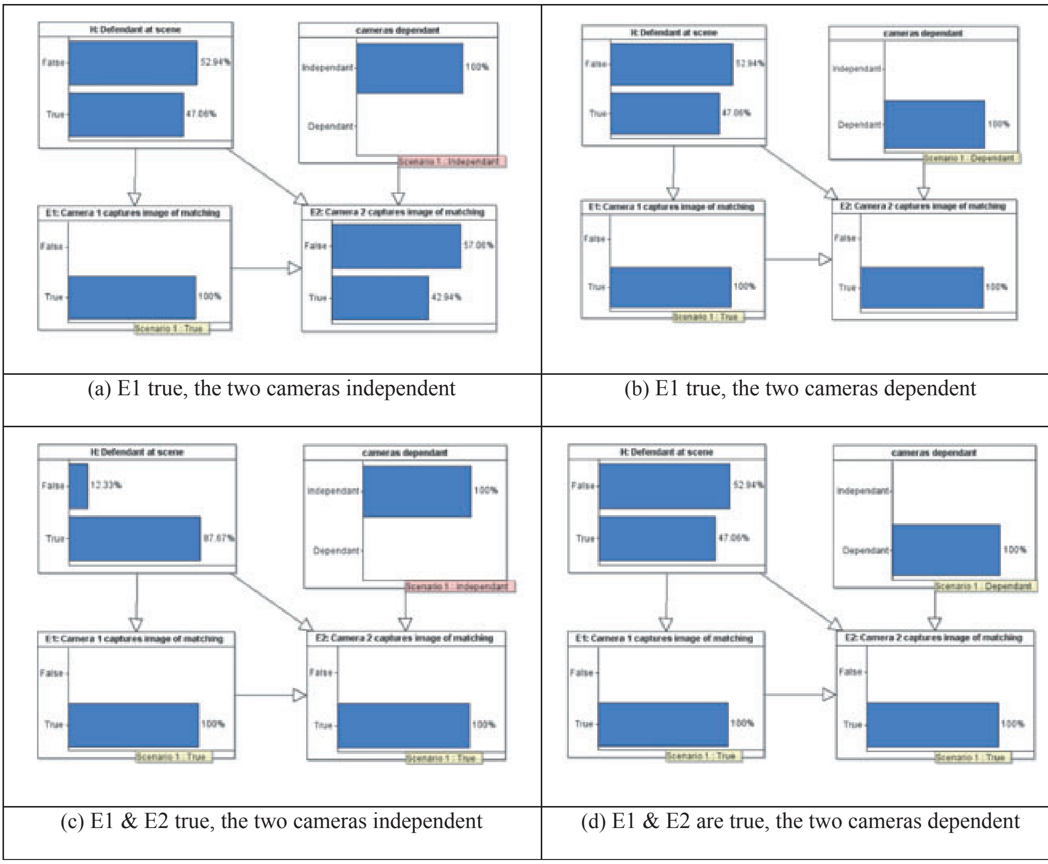

Fig. 15. (a) E1 is true, and the two cameras are independent. (b) E1 is true, and the two cameras are dependent. (c) E1 and E2 are true, and the two cameras are independent. (d) E1 and E2 are true, and the two cameras are dependent.

issue in the case of Sally Clark (Forrest, 2003). An important empirical question is the extent to which lay people (and legal professionals) are able to discount the value of dependent evidence. One empirical study (Schum & Martin, 1982) suggests that people sometimes ‘‘double-count’’ redundant evidence. This would lead to erroneous judgments, so it is vital to explore the generality of this error, and whether it can be alleviated by use of the BN framework proposed in this article.

There are other types of dependent evidence that require slightly different BN idioms that are beyond the scope of this study. These include (a) Dependent evidence through confirmation bias: In this case, there are two experts determining whether there is a forensic match (the type of forensics could even be different, such as DNA and fingerprinting). It has been shown (Dror & Charlton, 2006) that the second expert’s conclusion will be biased if he/she knows the conclusion of the first expert. (b) Dependent evidence through common biases, assumptions, and sources of inaccuracies. This is covered partly in Fenton and Neil (2012).

3.6. Alibi evidence idiom

A special case of additional direct dependency within the model occurs with so-called *alibi* evidence. In its most general form, alibi evidence is simply evidence that directly contradicts a prosecution hypothesis. The classic example of alibi evidence is an eyewitness statement contradicting the hypothesis that the defendant was present at the scene of the crime, normally by asserting that the defendant was in a different specific location. What makes this type of evidence special is that the hypothesis itself may directly influence the accuracy of the evidence such as when the eyewitness is either the defendant himself or herself or a person known to the defendant (see Lagnado, 2011). Fig. 16 shows the appropriate model with the revised dependency in the case where the witness is known to the defendant. A possible NPT for the node A1 (accuracy of alibi witness) is also shown in Fig. 16.

Imagine that the witness is the partner of the defendant. Then, what the NPT is saying is that if the defendant is not guilty, there is a very good chance the partner will provide an accurate alibi statement. However, if the defendant *is* guilty, there is a very good chance the partner's statement will be inaccurate. Of course, if the witness is an enemy of the defendant, the NPT will be somewhat inverted. But, with the NPT of Fig. 16, we can run the model and see the impact of the evidence in Fig. 17.

The model provides some very powerful analysis, notably in the case of conflicting evidence (i.e., where one piece of evidence supports the prosecution hypothesis and one piece supports the defense hypothesis). The following are the most interesting points to note:

1. When presented on their own ([b] and [c] respectively), both pieces of evidence lead to an increase in belief in their respective hypotheses. Hence, the alibi evidence leads to an increased belief in the defense hypothesis (not guilty) and the Closed Circuit Television (CCTV) evidence leads to an increased belief in the prosecution hypothesis (guilty). Obviously the latter is much stronger than the former because of the relative priors for accuracy, but nevertheless on their own they both provide support for their respective lawyers' arguments.

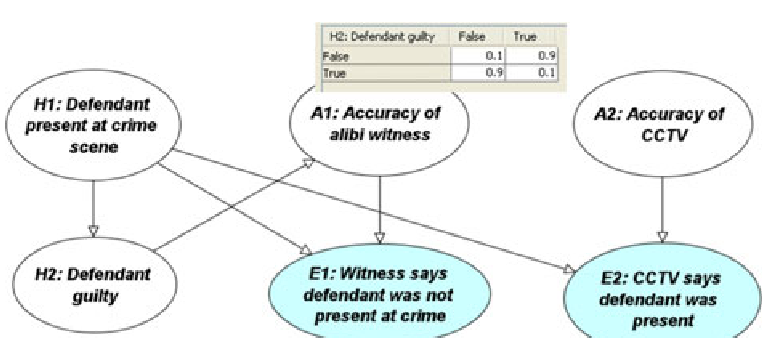

Fig. 16. Alibi evidence idiom, with Node Probability Table for A1 (accuracy of alibi witness).

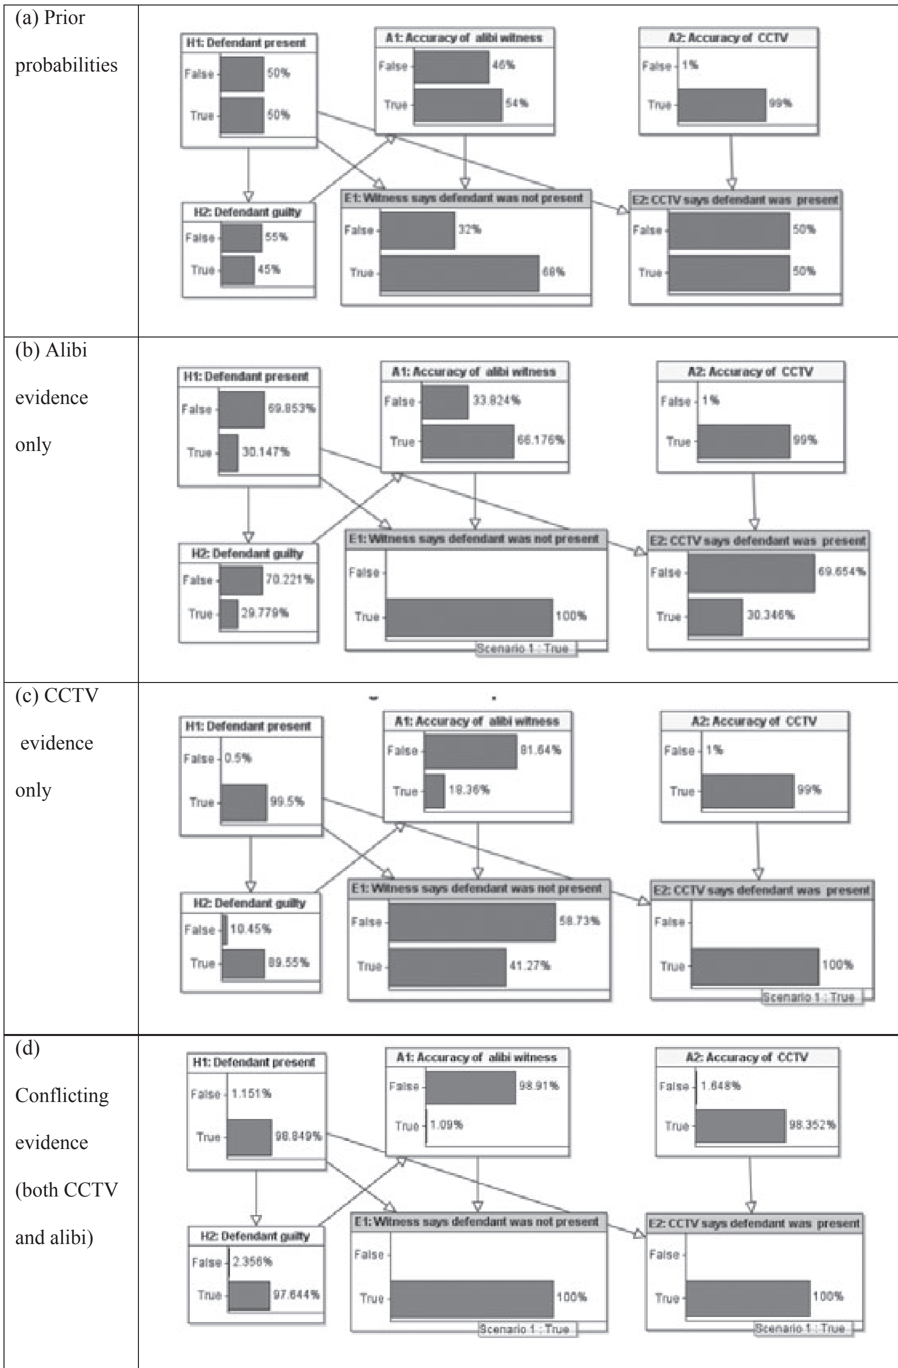

Fig. 17. Impact of alibi evidence. (a) Prior probabilities. (b) Alibi evidence only. (c) CCTV evidence only. (d) Conflicting evidence (both CCTV and alibi).

2. When both pieces of evidence are presented (d), we obviously have a case of conflicting evidence. If the pieces of evidence were genuinely independent, the net effect would be to decrease the impact of both pieces of evidence on their respective hypotheses compared with the single-evidence case. However, here because the alibi evidence is dependent on H2, the result is that the conflicting evidence actually strengthens the prosecution case even more than if the CCTV evidence was presented on its own. Specifically, because of the prior accuracy of the CCTV evidence, when this is presented together with the alibi evidence, it leads us to doubt the accuracy of the latter (we tend to believe the witness is lying) and hence, by backward inference, to increase the probability of guilt.

This analysis of alibi evidence has direct relevance to legal cases. Indeed, two key issues that arise when an alibi defense is presented in court are (a) whether the alibi provider is lying and (b) what inferences should be drawn if one believes that they are lying. In cases where an alibi defense is undermined, judges are required to give special instructions alerting the jury to the potential dangers of drawing an inference of guilt. In particular, the judge is supposed to tell the jury that they must be *sure* that the alibi provider has lied, and *sure* that the lie does not admit of an innocent explanation (Crown Court Benchbook, 2010). We maintain that the correct way to model alibi evidence, and to assess what inferences can be legitimately drawn from faulty alibis, is via the BN framework. Moreover, recent empirical studies show that ordinary people draw inferences in line with the proposed alibi idiom (Lagnado, 2011, 2012). For example, when given the scenario discussed above, judgments of the suspect's guilt are higher when both alibi evidence and disconfirming CCTV evidence are presented, than when CCTV evidence alone is presented. This holds true even though the alibi evidence by itself reduces guilt judgments. This pattern of inference is naturally explained by the supposition that the suspect is more likely to lie if he or she is guilty rather than innocent.

3.7. Explaining-away idiom

One of the most powerful features of BN reasoning is the concept of ''explaining away.'' An example of explaining away was seen in the evidence-accuracy idiom in Fig. 5a. The node E (evidence of blood match) has two parents H (defendant guilty) and A (accuracy of evidence) either of which can be considered being possible ''causes'' of E. Specifically, H being *true* can cause E to be true, and A being *false* can cause E to be true. When we know that the blood match evidence has been presented (i.e., E is true), then, as shown in Fig. 5c, the probability of *both* potential ''causes'' increases (the probability of H being true increases and the probability of A being false increases). Of the two possible causes, the model favors A being false as the most likely explanation for E being true. However, if we know for sure that A is true (i.e., the evidence is accurate), then, as shown in Fig. 5d, we have explained away the ''cause'' of E being true—it is now almost certain to be H being true. Hepler et al. (2007) consider ''explaining away'' as an explicit idiom as shown in Fig. 18a.

Hepler et al.'s example of their explaining-away idiom also turns out to be a special case of the evidence-accuracy idiom. In their example, the event is ''defendant confesses to the crime,'' and the causes are (1) defendant guilty and (2) defendant coerced by interrogating official. Using our terminology, the ''event'' is clearly a piece of evidence and cause 2 characterizes the accuracy of the evidence. However, it turns out that traditional ''explaining away'' does not work in a very important class of situations that are especially relevant for legal reasoning. These are the situations where the two causes are *mutually exclusive*, that is, if one of them is true then the other must be false. Suppose, for example, that we have evidence *E* that blood found on the defendant's shirt matches the victim's blood. As there is a small chance (let us assume 1%) that the defendant's blood is the same type as the victim's, there are two possible causes of this:

Cause 1: the blood on the shirt is from the victim.

Cause 2: the blood on the shirt is the defendant's own blood.

In this case, only one of the causes can be true. However, the standard approach to BN modeling will not produce the correct reasoning in this example (this issue is addressed

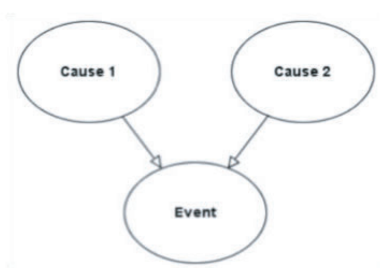

Blood on shirt is from victim	False		True	
Blood on shirt is from defendant	False	True	False	True
False	1.0	0.99	0.0	0.0
True	0.0	0.01	1.0	1.0

(a)

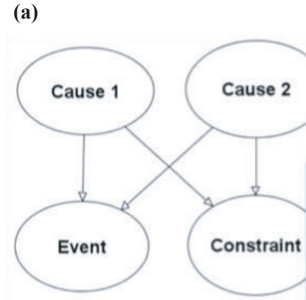

Cause 1	False		True	
Cause 2	False	True	False	True
Cause 1	0.0	1.0	0.0	0.0
Cause 2	0.0	0.0	1.0	0.0
Impossible	1.0	0.0	0.0	1.0

(b)

Fig. 18. (a) Explaining-away idiom, with possible Node Probability Table (NPT) for *E* (blood on shirt matches victim). (b) Explaining-away idiom with constraint node and its NPT.

indirectly in Pearl, 2011). Indeed, it is shown in Fenton, Neil, and Lagnado (2011) that in the general case of mutually exclusive causes, there is no way to correctly model the required behavior using a BN with the same set of nodes and states (even if we introduce dependencies between the cause nodes).

One solution would be to replace two separate cause nodes with a single cause node that has mutually exclusive states. Unfortunately, this approach is of little help in most legal reasoning situations because we will generally want to consider *distinct* parent and child nodes of the different causes, each representing distinct and separate causal pathways in the legal argument. For example, cause 2 may itself be caused by the defendant having cut himself in an accident; as cause 1 is not conditionally dependent on this event, it makes no sense to consider the proposition ''blood on shirt belongs to victim because the defendant cut himself in an accident.'' We cannot incorporate these separate pathways into the model in any reasonable way if cause 2 is not a separate node from cause 1. The solution described in Fenton et al. (2011) is to introduce a new node that operates as a constraint on the cause 1 and cause 2 nodes as shown in Fig. 18b.

As shown in the figure, the NPT of this new constraint node has three states: one for each causal pathway plus one for ''impossible.'' The impossible state can only be true when either (a) both causes 1 and 2 are false or (b) both causes 1 and 2 are true. Given our assumption that the causes are mutually exclusive, these conjoint events are by definition impossible, hence the choice of the third state label. To ensure that impossibility is excluded in the model operationally and to preserve the other prior probabilities, we enter what is called ''soft evidence'' on the constraint node according to the formula described in Fenton et al. (2011). This can be done using standard BN tools.

A complete example, which also shows how we can extend the use of the idiom to accommodate other types of deterministic constraints in the possible combinations of evidence/hypotheses, is shown in Fig. 19.

In this example, if we assume uniform priors for the nodes without parents, then once we set the evidence of the blood match as True, and the soft evidence on the Constraint node as describe above, we get the result shown in Fig. 20a.

In the absence of other evidence, this clearly points to cause 1 (blood on the shirt is from the victim) as being most likely and so strongly favors the guilty hypothesis. However, when we enter the evidence (a) victim and defendant have same blood type and (b) defendant has recent scar, then we get the very different result shown in Fig. 20b. This clearly points to cause 2 as being the most likely.

So, in summary, the key points about the above special ''explaining away'' idiom are as follows:

1. It should be used when there are two or more mutually exclusive causes of *E*, each with separate causal pathways.
2. The mutual exclusivity acts as a constraint on the state space of the model and can be modeled as a constraint.
3. When running the model, soft evidence must be entered for the constraint node to ensure that impossible states cannot be realized in the model.

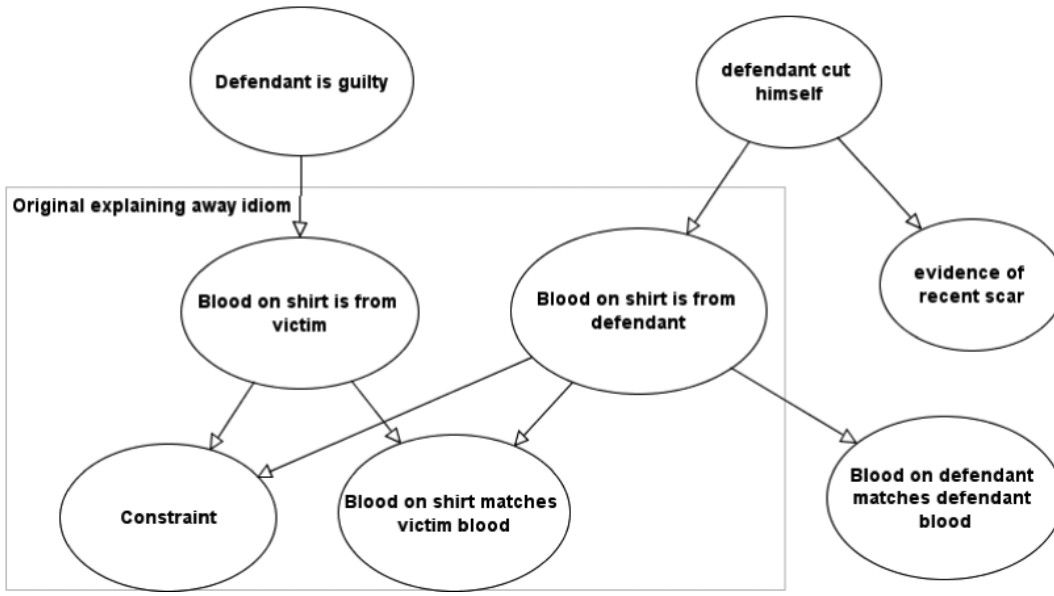

Fig. 19. Full example of mutually exclusive causes.

Using constraint nodes in this way also has the benefit of (a) revealing assumptions about the state space that would be otherwise tacit or implicit and (b) helping to keep causal pathways cleaner at a semantic level.

3.8. Direct versus circumstantial evidence

The idiom-based BN framework helps clarify the legal distinction between direct and circumstantial evidence (Roberts & Zuckerman, 2010). From the legal perspective, direct evidence is evidence that speaks directly to the issue to be proved, without any intermediate inferential step. For example, when a witness testifies to seeing the suspect commit the crime, or when the defendant confesses to the crime, this provides direct evidence of the guilt of the suspect. This kind of evidence can still be inconclusive—the witness might be unreliable for various reasons. Circumstantial evidence is indirect evidence; it speaks to the issue to be proved through an intermediate inferential step. For example, evidence of motive, opportunity, or identity is typical cases of circumstantial evidence. This kind of evidence is sometimes considered less probative than direct evidence; however, this depends greatly on the situation. Circumstantial evidence can be strong and convincing, especially in cases with forensic evidence such as DNA. Indeed some/many cases are decided entirely on circumstantial evidence. This distinction is readily mapped onto the BN framework. Direct evidence involves a single causal link from the issue to be proved to the evidence. If true, this evidence effectively proves the hypothesis in question. However, the evidence-accuracy idiom makes explicit that this evidence, although direct, might still be unreliable, and the unpacking of accuracy into

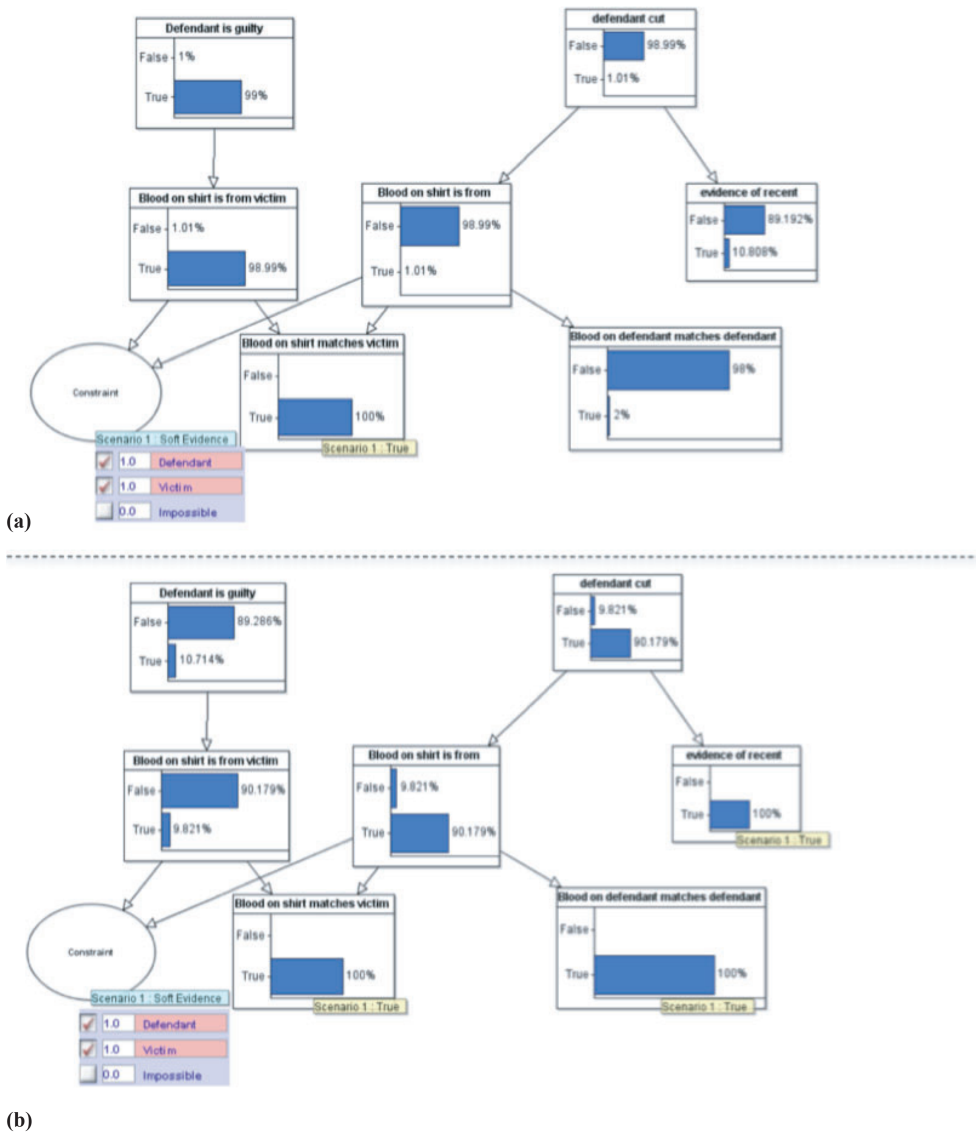

Fig. 20. (a) Evidence of blood match set to true. (b) New defense evidence is entered.

veracity, competence, and objectivity highlights possible reasons for this unreliability. Circumstantial evidence, in contrast, is linked to the issue to be proved via a causal path involving at least two steps, for instance, as evidence of motive or opportunity (as detailed in the corresponding idioms). The BN framework thus clarifies both that direct evidence can sometimes have greater probative value than circumstantial evidence, because it involves fewer inferential steps, but also that direct evidence can be severely weakened via the unreliability of its source.

4. Putting it all together: Vole example

Lagnado (2011) discussed the following fictional case based on Agatha Christie's play *Witness for the Prosecution* (Christie, 1953):

> Leonard Vole is charged with murdering a rich elderly lady, Miss French. He had befriended her, and visited her regularly at her home, including the night of her death. Miss French had recently changed her will, leaving Vole all her money. She died from a blow to the back of the head. There were various pieces of incriminating evidence: Vole was poor and looking for work; he had visited a travel agent to inquire about luxury cruises soon after Miss French had changed her will; the maid claimed that Vole was with Miss French shortly before she was killed; the murderer did not force entry into the house; Vole had blood stains on his cuffs that matched Miss French's blood type.
>
> As befits a good crime story, there were also several pieces of exonerating evidence: The maid admitted that she disliked Vole; the maid was previously the sole benefactor in Miss French's will; Vole's blood type was the same as Miss French's, and thus also matched the blood found on his cuffs; Vole claimed that he had cut his wrist slicing ham; Vole had a scar on his wrist to back this claim. There was one other critical piece of defence evidence: Vole's wife, Romaine, was to testify that Vole had returned home at 9.30 PM. This would place him far away from the crime scene at the time of Miss French's death. However, during the trial Romaine was called as a witness for the prosecution. Dramatically, she changed her story and testified that Vole had returned home at 10.10 PM, with blood on his cuffs, and had proclaimed: ''I've killed her.'' Just as the case looked hopeless for Vole, a mystery woman supplied the defence lawyer with a bundle of letters. Allegedly these were written by Romaine to her overseas lover (who was a communist!). In one letter she planned to fabricate her testimony in order to incriminate Vole, and rejoin her lover. This new evidence had a powerful impact on the judge and jury. The key witness for the prosecution was discredited, and Vole was acquitted.
>
> After the court case, Romaine revealed to the defence lawyer that she had forged the letters herself. There was no lover overseas. She reasoned that the jury would have dismissed a simple alibi from a devoted wife; instead, they could be swung by the striking discredit of the prosecution's key witness.

To model the case, Lagnado (2011) presented the causal model shown in Fig. 21. What we will now do is build the model from scratch using only the idioms introduced. In doing so, we demonstrate the effectiveness and simplicity of our proposed method (which provides a number of clarifications and improvements over the original model). Most important, we are able to run the model to demonstrate the changes in posterior guilt that result from presenting evidence in the order discussed in the example.

Step 1: Identify the key prosecution hypotheses (including opportunity and motive)

1. The ultimate hypothesis: ''H0: Vole guilty.''
2. Opportunity: ''Vole present.''
3. Motive: There are actually two possible motives, ''Vole poor'' and ''Vole in will.''

Step 2: Consider what evidence is available for each of the above and what is the accuracy of the evidence:

Evidence for H0. There is no direct evidence at all for H0 since no witness testifies to observing the murder. But what we have is evidence for are two hypotheses that depend on H0:

H1: Vole admits guilt to Romaine.

H2: Blood on Vole's shirt is from French.

Of course, neither of these hypotheses is guaranteed to be true if H0 is true, but this uncertainty is modeled in the respective NPTs.

The (prosecution) evidence to support H1 is the witness statement by Romaine. Note that Romaine's evidence of Vole's guilt makes her evidence of ''Vole present'' redundant (so there is no need for the link from ''Vole present'' to Romaine's testimony in the original model).

The issue of accuracy of evidence is especially important for Romaine's evidence. Because of her relationship with Vole, the H1 hypothesis influences her accuracy. Evidence to support H2 is that the blood matches French's. The evidence to support the opportunity ''Vole present'' is a witness statement from the Maid.

Step 3: Consider what defense evidence is available to challenge the above hypotheses.

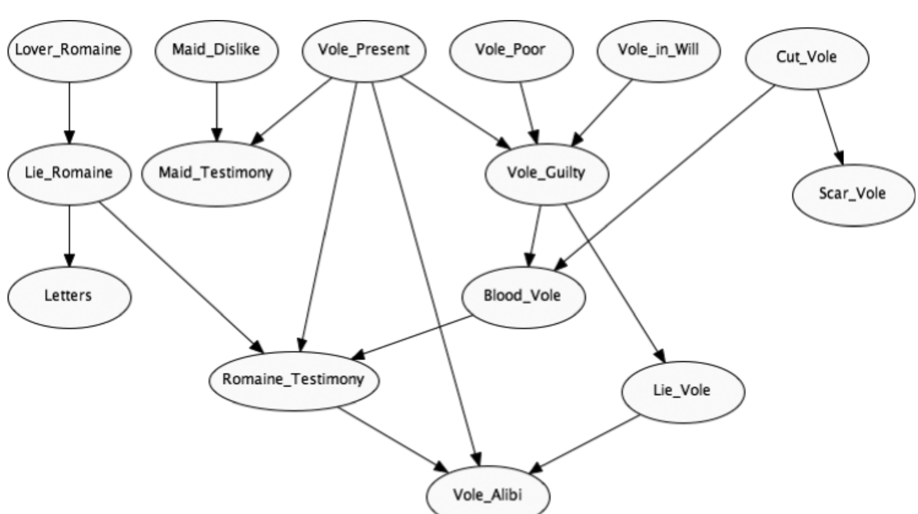

Fig. 21. Bayesian network model of ''Witness for the prosecution'' from Lagnado (2011).

The evidence to challenge H1 is the (eventual) presentation of the love letters and the introduction of a new (defense) hypothesis: ‘‘H4: Romaine has lover.’’

The evidence to challenge the opportunity ‘‘Vole present’’ is (a) to explicitly challenge the *accuracy* of the Maid's evidence and (b) Vole's own alibi evidence.

The evidence to challenge H2 is that the blood matches Vole's (i.e., Vole and French have the same blood type). Additionally, the defense provides an additional hypothesis: ‘‘H3: Blood on Vole is from previous cut’’ that depends on H2.

Finally, for simplicity, we shall assume that some evidence (such as the blood match evidence) is perfectly accurate and that the motives are stated (and accepted) without evidence. From this analysis, we get the BN shown in Fig. 22.

Note how the model is made up only from the idioms we have introduced (the blood match component is exactly the special ‘‘explaining-away’’ idiom example described above). With the exception of the node H5 (Romaine has lover), the priors for all parentless nodes are uniform. The node H5 has prior set to True = 10%. What matters when we run the model is not so much whether the probabilities are realistic but rather the way the model responds to evidence. Hence, Table 1 shows the effect (on probability of guilt) of the evidence as it is presented sequentially, starting with the prosecution evidence.

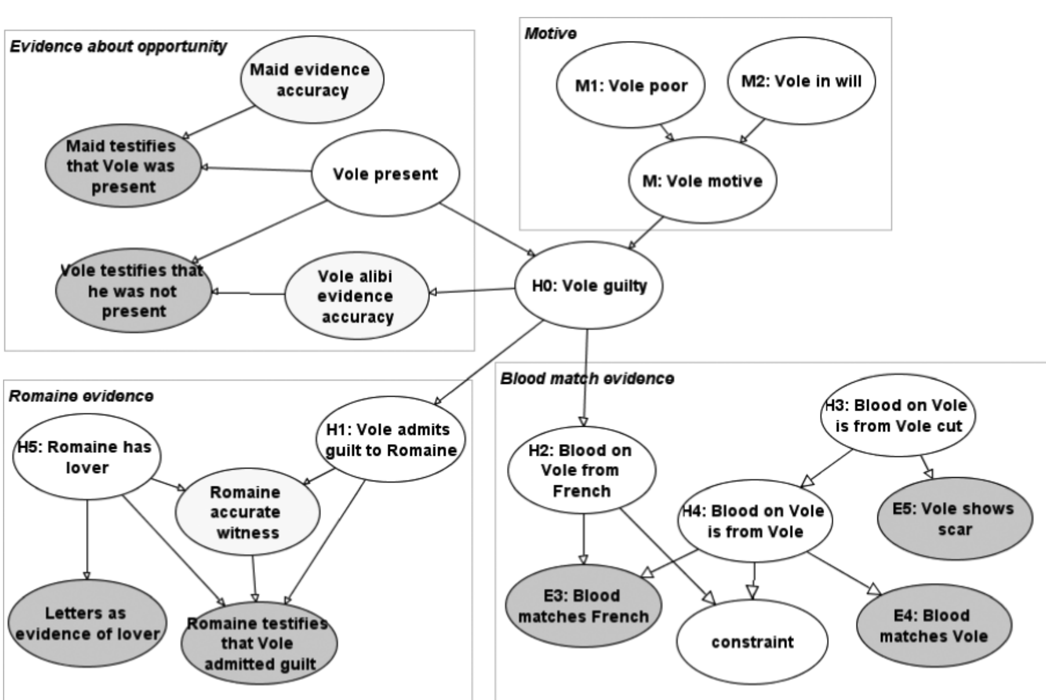

Fig. 22. Revised Bayesian network model of Vole case using idioms.

Table 1
Effect on probability of guilt of evidence presented sequentially in Vole case

Sequential Presentation of Evidence	H0 Vole Guilty Probability (as %)
Prosecution evidence presented	
1. Prior (no observations)	33.2
2. Motive evidence added (M1 and M2 = true)	35.8
3. Maid testifies Vole was present = true	52.6
4. E3 blood matches French evidence = true	86.5
5. Romaine testifies Vole admitted guilt = true	96.6
Defense evidence presented	
6. Vole testifies he was not present = true	96.9
7. Maid evidence accuracy = false	91.3
8. E4 Blood matches Vole = true	64.4
9. E5 Vole shows scar = true	40.4
10. Letters as evidence = true	14.9

The key points to note here are as follows:

1. The really ''big jump'' in belief in guilt comes from the introduction of the blood match evidence (at this point, it jumps from 52.6% to 86.5%). However, if Romaine's evidence had been presented before the blood match evidence, that jump would have been almost as great (52.6–81%).
2. Once all the prosecution evidence is presented (and bear in mind that at this point the defense evidence is set to the prior values), the probability of guilt seems overwhelming (96.6%).
3. If, as shown, the first piece of defense evidence is Vole's own testimony that he was not present, then the impact on guilt is negligible. This confirms that, especially when seen against stronger conflicting evidence, an alibi that is not ''independent'' is very weak. Although the model does not incorporate the intended (but never delivered) alibi statement by Romaine, it is easy to see that there would have been a similarly negligible effect, that is, Romaine's suspicions about the value of her evidence are borne out by the model.
4. The first big drop in probability of guilt comes with the introduction of the blood match evidence.
5. However, when all but the last piece of defense evidence is presented, the probability of Vole's guilt is still 40.4%—higher than the initial probability. Only when the final evidence—Romaine's letters—are presented do we get the dramatic drop to 14.9%. As this is considerably *less* than the prior (33.2%), this should certainly lead to a not guilty verdict if the jury were acting as rational Bayesians.

It is also worth noting the way the immediate impact of different pieces of evidence is very much determined by the order in which the evidence is presented. To emphasize this point, Fig. 23 presents a sensitivity analysis, in the form of a Tornado chart, of the impact of each possible piece of evidence *individually* on the probability of guilt. From this graph we can see, for example, that if all other observations are left in their prior state, the Vole blood

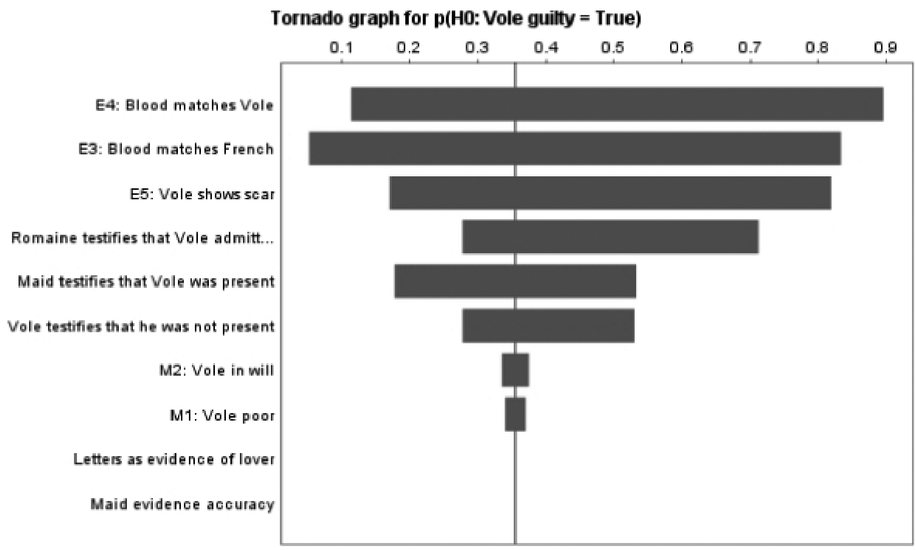

Fig. 23. Sensitivity analysis on guilty hypothesis.

match evidence has the largest impact on Vole guilty; when it is set to true, the probability of guilt drops to just over 10%, and when it is set to false, the probability of guilt jumps to nearly 90%. The French blood match evidence has a very similar impact. At the other extreme, the individual impact of the Romaine letters is almost negligible. This is because this piece of evidence only becomes important when the particular combination of other evidence has already been presented.

5. Road Map and conclusions

In this article, we have outlined a general framework for modeling legal arguments. The framework is based on Bayesian networks, but it introduces a small set of causal idioms tailored to the legal domain that can be reused and combined. This idiom-based approach allows us to model large bodies of interrelated evidence and capture inference patterns that recur in many legal contexts. The use of small-scale causal idioms fits well with the capabilities and constraints of human cognition, and thus it provides a practical method for the analysis of legal cases.

The proposed framework serves several complementary functions:

1. To provide a normative model for representing and drawing inferences from complex evidence, thus supporting the task of making rational inferences in legal contexts.
2. To suggest plausible cognitive models (e.g., representations and inference mechanisms) that explain how people manage to organize and interpret legal evidence.
3. To act as a standard by which to evaluate nonexpert reasoning (e.g., by jurors), where people depart from the rational model the BN approach provides methods and tools to improve judgments (especially with complex bodies of evidence).

The potential for the BN framework to illuminate the psychology of juror reasoning is relatively unexplored, but empirical findings thus far are encouraging. Research suggests that people naturally use causal models to organize and understand legal evidence (Pennington & Hastie, 1986, 1992) and can draw rational inferences in simple cases (Lagnado, 2011, 2012; Lagnado & Harvey, 2008). There is also growing evidence in the cognitive psychology literature that people's reasoning is often well captured within a general Bayesian framework (Griffiths, Kemp, & Tenenbaum, 2008; Griffiths & Tenenbaum, 2009; Hahn & Oaksford, 2007; Oaksford & Chater, 2007; Sloman, 2005). Of particular relevance is the work by Hahn, Oaksford, and colleagues on informal argumentation (Corner et al., 2011; Hahn & Oaksford, 2007; Hahn et al., 2009; Harris & Hahn, 2009; Jarvstad & Hahn, 2011). This body of research shows that people's evaluations of informal arguments fit within a Bayesian framework. It also shows, in line with the evidence-accuracy idiom proposed in this article, that people are sensitive to source reliability (Hahn et al., 2009; Jarvstad & Hahn, 2011), and that this can be modeled in Bayesian terms.

The current paper focuses on how BNs can capture legal arguments by representing the probabilistic causal relations between hypotheses and evidence. A different kind of approach to the formalization and visual representation of legal argument is provided by argumentation theory (Walton, 2008; Walton et al., 2008). This approach shares the main goals of analyzing and evaluating legal arguments, but it differs from the BN framework in several respects.

One key difference is that argumentation theory eschews the use of probability theory to handle the uncertainty inherent in legal arguments. Instead of probability, the concept of plausibility is introduced. However, this concept is not well defined, and the rules used to combine or propagate plausibilities lack a sound normative justification and conflict with everyday intuitions (see Hahn et al., in press for details). This is problematic for capturing complex legal argument, which requires the integration of large bodies of interrelated evidence. In contrast, the BN framework provides a coherent and well-defined system for combining and computing with probabilities. Another key difference lies in the inferential capacities of the two systems. BNs aim to model processes in the world, and it can be used to make hypothetical inferences and generate predictions about expected evidence. Moreover, the inferential mechanism (Bayesian updating) can sometimes lead to conclusions that were not apparent to the modeler (e.g., see page 20 of this paper). Argumentation diagrams (ADs), as presented by Walton and colleagues, do not aim to represent causal processes in the world, and thus do not generate novel inferences or predictions about what might have happened (or would have happened). Their main role is representational rather than inferential: They serve predominantly as aids to elucidate one's inferences rather than to generate them. In this sense, the two approaches are complementary, not contradictory.

A separate aim of argumentation theory is to model the process of legal dialog, including different types of burden of proof (Gordon & Walton, 2011; Walton, 2008). This is an important area of research that moves beyond the modeling of legal argument as construed in this article. Indeed, there are recent attempts to develop hybrid systems that combine argumentation theory with the BN approach (Bex, van Koppen, Prakken, & Verheij, 2010; Grabmair, Gordon, & Walton, 2010; Keppens, 2011). For example, Keppens (2011) examines the similarities and differences between BNs and ADs. He explains that although lacking the formal inference mechanisms of BNs, ADs enable richer and more diverse representations that make it suitable for marshaling all the information in a case in such a way that it is possible to identify relationships for evidential reasoning. As BNs and ADs offer different perspectives, they have the possibility to inform one another. Keppens' work extends the work by Hepler et al. (2007) in focusing on how the AD perspective could help inform the construction of BNs. This is a viewpoint we support and we see this work as highly complementary to what we propose in this paper. Keppens also proposes a method for extracting ADs from BNs.

Another related line of research is the use of model-based Bayesian methods for crime investigation (Keppens, Shen, & Price, 2011; Keppens & Zeleznikow, 2002, 2003). This work explores a different stage of the legal process, namely that of evidence collection, and uses model-based techniques for generating and analyzing plausible crime scenarios at the investigative phase. An important question for future research is the extent to which these different stages of the legal process—crime investigation versus interpretation of evidence in court—share common modeling techniques.

In sum, we believe that the probabilistic and inferential nature of BNs marks them out from argumentation theory and makes them an indispensible framework for legal arguments. In addition, the BN framework, together with the causal idiom approach, suggests numerous fresh avenues for empirical research. We have highlighted some of these throughout the paper; they include the explicit modeling of evidence accuracy, incorporating evidence of motive and opportunity, the inferences drawn from dependent evidence, the interpretation of alibi evidence, and the ''explaining away'' of competing causes. Many of these questions would not have been formulated without the introduction of a general framework for evidential reasoning. We are now in a position to systematically test how people reason in such contexts, and establish whether they conform to the prescripts of the BN models.

Irrespective of how these studies turn out, we believe that it is crucial to build a general inferential framework that allows us to understand and conduct legal argumentation. The proposed BN framework might prove a powerful guide to the representations and inferences that people actually use, but where people depart from this rational model, it will also afford a means for correcting these departures.

Acknowledgments

We would like to thank Adam Harris, Anne Hsu, and William Marsh (as well as the three reviewers) for helpful comments.

David Lagnado was supported by ESRC Grant RES- 062-33-0004, and Norman Fenton was supported by the EPSRC ImpactQM Programme.

References

AgenaRisk. (2012). Agena Ltd, Available at: http://www.agenarisk.com. Accessed September, 2012.

Aitken, C. G. G., T., Connolly, A., Gammerman, G., Zhang, D., Bailey, R., Gordon and R., Oldfield (1995). Bayesian belief networks with an application in specific case analysis. *Computational Learning and Probabilistic Reasoning*. A. Gammerman, John Wiley and Sons Ltd.

Aitken, C., & Taroni, F. (2004). *Statistics and the evaluation of evidence* (2nd ed.). Hoboken, NJ: Wiley.

Ashley, K. D. (1990). *Modeling legal argument: Reasoning with cases and hypotheticals*. Cambridge, MA: The MIT Press/Bradford Books.

Balding, D. J., & Donnelly, P. (1994). Prosecutor's fallacy and DNA evidence. *Criminal Law Review*, *October*, 711–722.

Bankowski, Z., White, I., & Hahn, U. (1995). *Informatics and the foundations of legal reasoning*. Dordrecht the Netherlands: Kluwer Academic Publishers.

Bex, F., van Koppen, P., Prakken, H., & Verheij, B. (2010). A hybrid formal theory of arguments, stories and criminal evidence. *Artificial Intelligence and Law*, *18*, 123–152.

Bovens, L., & Hartmann, S. (2003). *Bayesian epistemology*. Oxford, England: Oxford University Press.

Casscells, W., Schoenberger, A., & Graboys, T. B. (1978). Interpretation by physicians of clinical laboratory results. *New England Journal of Medicine*, *299*, 999–1001.

Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, *4*, 55–81.

Christie, A. (1953/2002). *Witness for the prosecution*. London: Harper Collins Publishers Ltd.

Corner, A., Hahn, U., & Oaksford, M. (2011). The psychological mechanism of the slippery slope argument. *Journal of Memory and Language*, *64*, 133–152.

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavorial & Brain Sciences.*, *24*, 87–185.

Crown Court Benchbook (2010). Judicial Studies Board.

Dawid, A. P. (2002). Bayes's theorem and weighing evidence by juries. In R. Swinburne (Ed.), *Bayes's theorem* (pp. 71–90). Proceedings of the British Academy. Oxford, England: Oxford University Press.

Dawid, A. P., & Evett, I. W. (1997). Using a graphical model to assist the evaluation of complicated patterns of evidence. *Journal of Forensic Sciences*, *42*, 226–231.

Dror, I. E., & Charlton, D. (2006). Why experts make errors. *Journal of Forensic Identification*, *56*(4), 600–616.

Dror, I. E., & Hampikian, G. (2011). Subjectivity and bias in forensic DNA mixture interpretation. *Science & Justice*, *51(4)*, 204–208.

Edwards, W. (1991). Influence diagrams, Bayesian imperialism, and the collins case: An appeal to reason. *Cardozo Law Review*, *13*, 1025–1079.

Ericsson, K., Charness, N., Feltovich, P., & Hoffman, R. (2006). *Cambridge handbook of expertise and expert performance*. Cambridge, UK: Cambridge University Press.

Ericsson, K., & Kintsch, W. (1995). Long-term working memory. *Psychological Review*, *102*, 211–245.

Evett, I. W., & Weir, B. S. (1998). *Interpreting DNA evidence: Statistical genetics for forensic scientists*. Sunderland, MA: Sinauer Associates Inc.

Faigman, D. L., & Baglioni, A. J. (1988). Bayes' theorem in the trial process. *Law and Human Behavior*, *12*, 1–17.

Fenton, N. E. (2011). Science and law: Improve statistics in court. *Nature*, *479*, 36–37.

Fenton, N., & Neil, M. (2010). Comparing risks of alternative medical diagnosis using Bayesian arguments. *Journal of Biomedical Informatics*, *43*, 485–495.

Fenton, N., & Neil, M. (2011). Avoiding legal fallacies in practice using Bayesian networks. *Australian Journal of Legal Philosophy*, *36*, 114–150.

Fenton, N. E., & Neil, M. (2012). On limiting the use of Bayes in presenting forensic evidence, Extended draft, Available at: https://www.eecs.qmul.ac.uk/~norman/papers/likelihood_ratio.pdf. Accessed 24 September 2012.

Fenton, N., Neil, M., & Lagnado, D., (2011). Modelling mutually exclusive causes in Bayesian networks, Extended draft, Available at: http://www.eecs.qmul.ac.uk/~norman/papers/mutual_IEEE_format_version.pdf. Accessed 24 September 2012.

Fienberg, S. E., & Schervish, M. J. (1986). Relevance of Bayesian inference for the presentation of statistical evidence and for legal decisionmaking, the symposium. *Boston University Law Review*, *66*, 771–798.

Finkelstein, M. O., & Levin, B. A. (2001). *Statistics for lawyers*. New York: Springer.

Forrest, A. R. (2003). Sally Clark – A lesson for us all. *Science & Justice*, *43*, 63–64.

Friedman, R. (1987). Route analysis of credibility and hearsay. *The Yale Law Journal*, *96*(4), 668–742.

Gobet, F., Lane, P., Croker, S., Cheng, P., Jones, G., Oliver, I., & Pine, J. (2001). Chunking mechanisms in human learning. *Trends in Cognitive Science*, *5*, 236–243.

Good, P. I. (2001). *Applying statistics in the courtroom: A new approach for attorneys and expert witnesses*. Boca Raton, FL: CRC Press.

Gordon, T. F., & Walton, D. (2011). A formal model of legal proof standards and burdens. In F. van Emeren, B. Garssen, J. A. Blair, & G. R. Mitchell (Eds.), *7th Conference on Argumentation of the International Society for the Study of Argumentation (ISSA 2010)* (pp. 644–655). Amsterdam: Sic Sac.

Grabmair, M., Gordon, T. F., & Walton, D. (2010). Probabilistic semantics for the carneades argument model using Bayesian networks. In P. Baroni, F. Cerutti, M. Giacomin, & G. R. Simari (Eds.), *Computational Models of Argument: Proceedings of COMMA 2010* (pp. 255–266). Amsterdam: IOS Press.

Griffiths, T. L., Kemp, C., & Tenenbaum, J. B. (2008). Bayesian models of cognition. In R. Sun (Ed.), *The Cambridge handbook of computational cognitive modelling* (pp. 59–100). Cambridge, England: Cambridge University Press.

Griffiths, T. L., & Tenenbaum, J. B. (2009). Theory-based causal induction. *Psychological Review*, *116*, 661–716.

Hahn, U., Harris, A. J. L., & Corner, A. (2009). Argument content and argument source: An exploration. *Informal Logic*, *29*, 337–367.

Hahn, U., & Oaksford, M. (2007). The rationality of informal argumentation: A Bayesian approach to reasoning fallacies. *Psychological Review*, *114*, 704–732.

Hahn, U., Oaksford, M., & Harris, A. J. L. (in press). Testimony and argument: A Bayesian perspective. In F. Zenker (Ed.), *Bayesian argumentation*. New York: Springer Library.

Halford, G. S., Cowan, N., & Andrews, G. (2007). Separating cognitive capacity from knowledge: A new hypothesis. *Trends in Cognitive Science*, *11*, 236–242.

Harris, A. J. L., & Hahn, U. (2009). Bayesian rationality in evaluating multiple testimonies: Incorporating the role of coherence. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*, 1366–1372.

Hepler, A. B., Dawid, A. P., & Leucari, V. (2007). Object-oriented graphical representations of complex patterns of evidence. *Law, Probability & Risk*, *6*, 275–293.

Hugin, A. S. (2011). Hugin Expert. Available at: http://www.hugin.com. Accessed 24 September 2012.

Huygen, P. E. M. (2002). *Use of Bayesian Belief Networks in legal reasoning. In 17th BILETA Annual Conference*. Amsterdam: Free University.

Jackson, G., Jones, S., Booth, G., Champod, C., & Evett, I. W. (2006). The nature of forensic science opinion—a possible framework to guide thinking and practice in investigations and in court proceedings. *Science and Justice*, *46*, 33–44.

Jarvstad, A., & Hahn, U. (2011). Source reliability and the conjunction fallacy. *Cognitive Science*, *35*, 682–711.

Jowett, C. (2001). Lies, damned lies, and DNA statistics: DNA match testing. Bayes' theorem, and the criminal courts. *Medical Science & Law*, *41*, 194–205.

Kadane, J. B., & Schum, D. A. (1996). *A probabilistic analysis of the Sacco and Vanzetti evidence*. New York: John Wiley.

Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and biases*. New York: Cambridge University Press.

Keppens, J. (2011). On extracting arguments from Bayesian network representations of evidential reasoning. In Proceedings of the 13th International Conference on Artificial Intelligence and Law (pp. 141–150).

Keppens, J., Shen, Q., & Price, C. (2011). Compositional Bayesian modelling for computation of evidence collection strategies. *Applied Intelligence*, *35*, 134–161.

Keppens, J., & Zeleznikow, J. (2002). On the role of model-based reasoning in decision support in crime investigation. In Proceedings of the 3rd International Conference on Law and Technology (pp. 77–83).

Keppens, J., & Zeleznikow, J. (2003). A model based reasoning approach for generating plausible crime scenarios from evidence. In Proceedings of the 9th International Conference on Artificial Intelligence and Law (pp. 51–59).

Krynski, T. R., & Tenenbaum, J. B. (2007). The role of causality in judgment under uncertainty. *Journal of Experimental Psychology: General*, *136*, 430–450.

Lagnado, D. (2011). Thinking about evidence. In P. Dawid, W. Twining, & M. Vasilaki (Eds.), *Evidence, inference and enquiry. Proceedings of the British Academy/171*, Oxford: Oxford University Press.

Lagnado, D. A., Fenton, N., & Neil, M. (2012). Legal idioms: a framework for evidential reasoning. *Argument and Computation* 1–18.

Lagnado, D. A., & Harvey, N. (2008). The impact of discredited evidence. *Psychonomic Bulletin & Review*, *15*, 1166–1173.

Lagnado, D. A., Waldmann, M. A., Hagmayer, Y., & Sloman, S. A. (2007). Beyond covariation. Cues to causal structure. In A. Gopnik & L. E. Schultz (Eds.), *Causal learning: Psychology, philosophy, and computation* (pp. 154–172). New York: Oxford University Press.

Matthews, R. (1997). Tipping the scales of justice. *New Scientist*, *2112*, 18–19.

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, *63*, 81–97.

Neil, M., Fenton, N., & Nielsen, L. (2000). Building large-scale Bayesian Networks. *Knowledge Engineering Review*, *15*, 257–284.

Neil, M., Tailor, M., & Marquez, D. (2007). Inference in Bayesian networks using dynamic discretisation. *Statistics and Computing*, *17(3)*, 219–233, 2007.

Oaksford, M., & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford, England: Oxford University Press.

Oaksford, M., & Chater, N. (Eds.) (2010). *Cognition and conditionals: Probability and logic in human thinking*. Oxford, England: Oxford University Press.

Pearl, J. (1988). *Probabilistic reasoning in intelligent systems*. Palo Alto, CA: Morgan Kaufmann.

Pearl, J. (2011). The algorithmization of counterfactuals. *Annals for Mathematics and Artificial Intelligence*, *61*(1), 29–39.

Pennington, N., & Hastie, R. (1986). Evidence evaluation in complex decision making. *Journal of Personality and Social Psychology*, *51*, 242–258.

Pennington, N., & Hastie, R. (1992). Explaining the evidence: Test of the story model for juror decision making. *Journal of Personality and Social Psychology*, *62*, 189–206.

Prakken, H. (1997). Logical Tools for Modelling Legal Argument: A Study of Defeasible Reasoning in Law, *Law and Philosophy Library*, *32*, Dordrecht: Kluwer Boston/London.

R. v. George (2007). EWCA Crim 2722.

Redmayne, M. (1995). DNA evidence, probability and the courts. *Criminal Law Review*, *464*, 464–482.

Roberts, P., & Zuckerman, A. (2010). *Criminal evidence* (2nd ed.). Oxford, England: Oxford University Press.

Robertson, B., & Vignaux, T. (1995). *Interpreting evidence: Evaluating forensic science in the courtroom*. Chichester: John Wiley.

Robertson, B., & Vignaux, T. (1997). Bayes' theorem in the court of appeal. *The Criminal Lawyer*, January, 4–5.

Schum, D. A. (2001). *The evidential foundations of probabilistic reasoning*. Evanston, IL: Northwestern University Press.

Schum, D. A., & Martin, A. W. (1982). Formal and empirical research on cascaded inference in jurisprudence. *Law & Society Review*, *17*, 105–151.

Sloman, S. A. (2005). *Causal models*. Cambridge, MA: Oxford University Press.

Sloman, S. A., & Lagnado, D. A. (2005). Do we do? *Cognitive Science*, *29*, 5–39.

Steyvers, M., & Griffiths, T. L. (2008). Rational analysis as a link between human memory and information retrieval. In N. Chater & M. Oaksford (Eds.), *The probabilistic mind: Prospects from rational models of cognition* (pp. 327–347). Oxford: Oxford University Press.

Taroni, F., Aitken, C., Garbolino, P., & Biedermann, A. (2006). *Bayesian networks and probabilistic inference in forensic science*. Chichester: John Wiley.

Thompson, W. C. (2009). Painting the target around the matching profile: The Texas sharpshooter fallacy in forensic DNA interpretation. *Law, Probability and Risk*, *8*(3), 257–276.

Tillers, P. (2011). Trial by Mathematics – Reconsidered. Cardozo Legal Studies Research Paper No. 333.

Tillers, P., & Green, E. (1988). *Probability and inference in the law of evidence the uses and limits of Bayesianism*. Dordrecht The Netherlands: Kluwer Academic Publishers.

Walton, D. (2008). *Witness testimony evidence: Argumentation, artificial intelligence, and law*. Cambridge, UK: Cambridge University Press.

Walton, D., Reed, C., & Macagno, F. (2008). *Argumentation schemes*. Cambridge, England: Cambridge University Press.

Wigmore, J. H. (1913). The problem of proof. *Illinois Law Journal*, *8*(2), 77–103.

Zukerman, I. (2010). Arguing with a Bayesian intelligence. In R. Dechter, H. Geffner, & J. Y. Halpern (Eds.), *Heuristics, probability and causality: A tribute to Judea Pearl* (pp. 271–292). London: College Publications.

Zukerman, I., & George, S. (2005). A probabilistic approach for argument interpretation. *User Modelling and User-Adapted Interaction*, *15*(1–2), 5–53.