De-biasing role induced bias using Bayesian networks

MARK SCHWEIZER

The merits of using subjective probability theory as a normative standard for evidence evaluation by legal fact-finders have been hotly debated for decades. Critics argue that formal mathematical models only lead to an apparent precision that obfuscates the ad-hoc nature of the many assumptions that underlie the model. Proponents of using subjective probability theory as normative standard for legal decision makers, specifically proponents of using Bayesian networks as decision aids in complex evaluations of evidence, must show that formal models have tangible benefits over the more natural, holistic assessment of evidence by explanatory coherence. This paper demonstrates that the assessment of evidence using a Bayesian network parametrized with values obtained from the decision makers greatly reduces role-induced bias, a bias that has been largely resistant to de-biasing attempts so far. This shows that using Bayesian networks as decision aids can benefit legal decision making.

I. Introduction

For over 40 years it has been proposed that forensic evidence should be assessed according to the rules of subjective probability theory, with the fact finder updating his or her probability of the relevant event depending on the strength of the evidence received (Lempert 1977; Kaye 1979; Schum 1988; Tillers 1988). It has been shown that most people violate the norms of subjective probability theory when updating their beliefs (Kahneman und Tversky 1973; Tversky und Kahneman 1982, 1983). Ensuring the coherence of partial beliefs, in the sense of subjective probability theory, quickly becomes computationally intractable (Callen 1982). The latter problem is solved by so called "Bayesian networks", which allow the compact representation of the full joint probability distribution using a directed graph and conditional probabilities (Pearl 1988). A number of authors have proposed using Bayesian inference networks for the evaluation of evidence in legal contexts (Edwards 1991; Robertson und Vignaux 1992; Kadane und Schum 1996; Taroni et al. 2006; Juchli et al. 2012; Lagnado et al. 2012; Fenton et al. 2013), some specifically suggesting that Bayesian networks as decision aids may help overcoming the shortcomings of human probabilistic reasoning (Fenton und Neil 2011).

Critics of the use of subjective probability theory as normative standard for the evaluation of evidence in legal cases have argued that mathematical calculations lead to a mere appearance of precision that is deceptive (Tribe 1971). Formal modelling required the assessment of many probabilities for which there were simply no bases in facts, burdening the fact finder with a formal exercise without a redeeming benefit (Thagard 2003). The more natural, holistic assessment of evidence based on explanatory coherence was therefore not only less burdensome, but superior and to be preferred from a normative point of view (Allen 1991; Thagard 2003; Amaya 2008).

Ultimately, the debate whether a holistic assessment of the evidence based on coherent reasoning or an assessment with the aid of a Bayesian network is to be normatively preferred can only be settled if it could be shown that one or the other approach leads to a representation of the facts that more closely reflects the ground truth (if one subscribes to a notion of "ground truth", i.e., an external state of the world that is independent of an observer, at all. If not, the entire fact-finding exercise seems futile, anyway). It will be nearly impossible to ever ascertain the superiority of one of the approaches in this sense, however, because the ground truth is unknown in most cases. The best we can hope for is showing that an approach leads to an elimination or at least reduction in biases and fallacies of reasoning well established by laboratory experiments.

On such well-documented bias is the dramatic effects of social roles on behaviour (Janis und King 1954; Zimbardo 1965; Vidmar und Laird 1983; Thompson und Loewenstein 1992). People often act in accordance with their role. Role-induced bias is not purely motivational, it has a cognitive component that subjects are unaware of (Engel und Glöckner 2013). It has been argued that this bias is highly relevant for law because roles influence choices (Sunstein, 1996), as well as preparatory acts, such as the search for evidence (O'Brien 2009). Role-induced bias might also be one of the causes for attorneys' overconfidence in predicting case outcomes (Goodman-Delahunty et al. 2010). Attempts at de-biasing role induced bias have met with limited success (Simon 2004; O'Brien 2009; Engel und Glöckner 2013).

The contribution of this paper lies in the experimental demonstration that assessing evidence with the aid of a "Bayesian network" significantly reduces the well-known role-induced bias, lending support to the proposition that legal fact finders should rely on Bayesian networks as decision aids. The rest of this paper is structured as follows: first, an empirically well-grounded descriptive theory of evidence evaluation is introduced, and it is shown how it leads to role-induced bias. Secondly, the debate on the merits of using subjective probability theory as a normative standard for evidence evaluation is briefly touched upon, and it is argued that the case for Bayesian networks as decision aids can only be made if it can be shown that they are useful for avoiding common biases. Thirdly, the hypothesis, method and results of the experiment are explained. A discussion of the results follows and leads to a conclusion.

II. Coherent reasoning in evidence evaluation

A. Coherence construction by parallel constraint satisfaction

Models of coherent reasoning assume that fact finders base their decision on constructing and evaluating coherent interpretations or stories from the available pieces of evidence (Pennington und Hastie 1988). Parallel constraint satisfaction (PCS) models using symbolic networks are a more formal evolution of the classic story model of evidence evaluation (Simon et al. 2004; Glöckner und Betsch 2008; Holyoak und Simon 1999; Engel und Glöckner 2013; Byrne 1995). PCS models assume that automatically spreading activation processes leads to constructing the best (i.e., most coherent) interpretation under parallel consideration of all constraints resulting from the evidence, the fact finders background knowledge and all logical relations (Thagard 1989; Thagard und Verbeurgt 1998). PCS problems can be represented in a connectionist network, which allows their approximative solution (Read et al. 1997). In a connectionist network, positively linked variables excite each other while negatively linked variables inhibit each other. In an iterative process, activation spreads through the network. The core feature of constraint satisfaction mechanisms is that the connectionist network will reconfigure itself until the constraints settle at a point of maximal coherence. This process forces coherence upon even an initially incoherent mental representation. Since the links between nodes in a connectionist network are bidirectional, the evidence influences the hypotheses, but the activation of the hypotheses also influences the interpretation of the evidence (Holyoak und Simon 1999). The formation of coherence in an iterative process therefore leads to a *polarization of the evidence*: evidence that supports the emerging decision is strongly endorsed while contradicting evidence is dismissed, rejected, or ignored (Simon 2004; DeKay 2015). This systematic revaluation of the evidence is called a coherence shift (Holyoak und Simon 1999) or coherence effect (Simon 2004). Coherence shifts have been demonstrated in a wide variety of tasks (Brownstein 2003) and particularly for legal judgments (Holyoak und Simon 1999; Carlson und Russo 2001; Hope et al. 2004; Lundberg 2004; Simon 2004; Engel und Glöckner 2013). Coherence shifts occur unconsciously and affect evidence that is logically unrelated (Holyoak und Simon 1999).

B. Emergence of role-induced bias

Role-induced bias can be caused by deliberate motivated reasoning (Kunda, 1990). Individuals might come to the conclusion which it is mandated by their role, in line with research in social psychology (Janis und King 1954; Zimbardo 1965). Yet the bias might also emerge unintentionally (Kunda und Thagard 1996; Monroe und Read 2008; Engel und Glöckner 2013), which can be explained by the automatic activation of unconscious goals in PCS networks. If the bias was intentionally or unintentionally formed by motivated reasoning, one should expect it to disappear when motivational goals are changed. According to PCS models, this is not the case. Once induced, the bias prevails even if goals change, because interpretations, once they have been formed, stabilize themselves by coherence shifts in the respective direction (Read et al. 1997). After the preferred interpretation is formed, all pieces of evidence are viewed in the light of this interpretation, and their evaluation is biased to support it (DeKay 2015). New information will be re-evaluated to match the overall interpretation (Simon 2004).

Engel und Glöckner 2012 show that role-induced bias is persistent and largely the result of unconscious thinking. Controlling for differences in information search, role-induced bias persists even when subjects have a high monetary incentive – EUR 100 – to accurately predict the outcome. Subjects in the role of prosecutor consider the probability of a guilty verdict significantly higher than subjects in the role of defense. The difference is mediated by coherence shifts (Engel und Glöckner 2012).

As predicted by PCS models, role-induced bias is more than just motivated reasoning or selective information search. Once "trapped" in an interpretation, it is hard to leave it and to come to a different interpretation.

Coherence shifts modify the interpretation of information and stabilize these interpretations once they have been formed (Engel und Glöckner 2013). This aligns well with older research showing that once a person has adopted a (social) theory, he or she is reluctant to give it up even if the theory has been discredited (Anderson et al. 1980). It could also explain why attempts at de-biasing role-induced bias have been met with limited success (Simon 2004; O'Brien 2009). The pure assignment of a role, even if there is no self-serving element, may have behavioural effects that cannot easily be reversed.

III. Subjective probability theory as normative standard of evidence evaluation

A. Subjective probability theory in legal reasoning

According to the subjective interpretation of probability, probability is a degree of belief (de Finetti 1937). Unlike the frequentist interpretation, the subjective interpretation of probability allows to speak intelligibly of the "probability" of a particular case. "Subjectivists" or "Bayesians" believe that the partial beliefs of a subject should (normatively) not violate the axioms of probability theory. A variety of arguments can be made why degrees of belief should conform to the axioms of probability theory. The least technical one is that unless the beliefs of a subject conform to the axioms of probability theory, the subject can be made the victim of a "Dutch book", a set of bets that incurs him a certain loss, no matter how the state of the world turns out (de Finetti 1937; Christensen 2007). From the axioms of probability theory follows immediately that the conditional probability of A given B is calculated according to Bayes theorem. For adherents of subjective probability theory, a subject should update her prior belief in the truth of a proposition A upon learning new information B according to Bayes theorem (Good 1950).

Whether subjective probability theory is a useful model for the formation of a belief in the context of the forensic evaluation of evidence is subject to a debate that has been likened to a 40 year war (Park et al. 2010; Aitken 2012). Some people take issue with the betting paradigm of subjective probability theory (Cohen 1977), while others are convinced that the expression of degrees of belief that are not grounded in observed relative frequencies in mathematical terms will lead to "wholly inaccurate, and misleadingly precise, conclusions" (Tribe 1971). This is essentially also the position of the Appeal Court of England and Wales (*R v. T* [2010], EWCA Crim 2439; see also *Nulty and others v. Milton Keynes Borough Council*, [2013] EWCA Civ 15) and the German Federal Court of Justice (NJW 1989, 3161).

B. Bayesian networks as decision aids in legal fact finding

Irrespective of the merits of subjective probability theory for the evaluation of evidence in text book examples, the problem remains that the "pen and paper" approach fails in the face of more complex evidential assessments. Ensuring the coherence of partial beliefs in the sense of subjective probability theory quickly becomes impossible without some sort of decision aid (Charniak 1991). Bayesian networks, also referred to as "belief nets" (Darwiche 2009), are such a decision aid. A number of authors have suggested using Bayesian inference networks for the evaluation of evidence in legal contexts (Edwards 1991; Robertson und Vignaux

1992; Kadane und Schum 1996; Taroni et al. 2006; Juchli et al. 2012; Lagnado et al. 2012; Fenton et al. 2013; Zoete et al. 2014; Fenton et al. 2016), some specifically suggesting that Bayesian networks as decision aids may be able to overcome the shortcomings of human probabilistic reasoning (Fenton und Neil 2011).

A Bayesian network is a directed acyclic graph in which a node (variable) is connected by a directed edge to another node if the variable represented by the node has a direct influence on the other variable (for a general introduction into Bayesian networks see (Taroni et al. 2006), p. 33 seq.). A conditional probability table is associated with each node (root nodes are only associated with "unconditional" or "prior" probabilities), which gives the probability for each mutually exclusive state of the variable given its parents (a parent of a node is an immediate ancestor of this node, i.e., any node that is *directly* connected to the node). In the network used in this paper, each variable can take on only two states which can be interpreted as "true" and "false". Using the concept of conditional independence, Bayesian networks can represent all direct *and indirect* dependencies of the problem domain by only explicitly showing the *direct* dependencies.

The following simple example, adapted from (Taroni et al. 2006) p. 39, may illustrate the concept. The subject holds some prior beliefs about the fairness of a coin, which can be fair (heads and tails on opposite sides), tails only or heads only. H is the variable that represents this prior belief, and H = fair, H = only heads and H = only tails are the three mutually exclusive states it can take. The subject now observes the outcome of a first throw of this coin (she cannot examine the coin). The variable E_1 represents this evidence (while E_2, E_3 stand for the second and third toss), and it can take the states E_1 = heads or E_1 = tails. This obviously tells the subject something about the fairness of the coin, and this in turn will influence her expectation that the next toss of the coin lands on heads (see Figure 1 a)). However, *given* that the coin is in any of its states, the outcome of the first toss will not tell the subject anything about the further outcomes. The variables E_1, E_2 and E_3 are *conditionally independent* given H. This knowledge of the conditional independencies is brought to the table by the human expert who knows that the first toss of a fair coin tells her nothing about the probable outcome of the second toss and allows a more parsimonious representation of the problem as given in Figure 1 b). Figure 1 b) also shows the (conditional) probability tables associated with each node of the network.

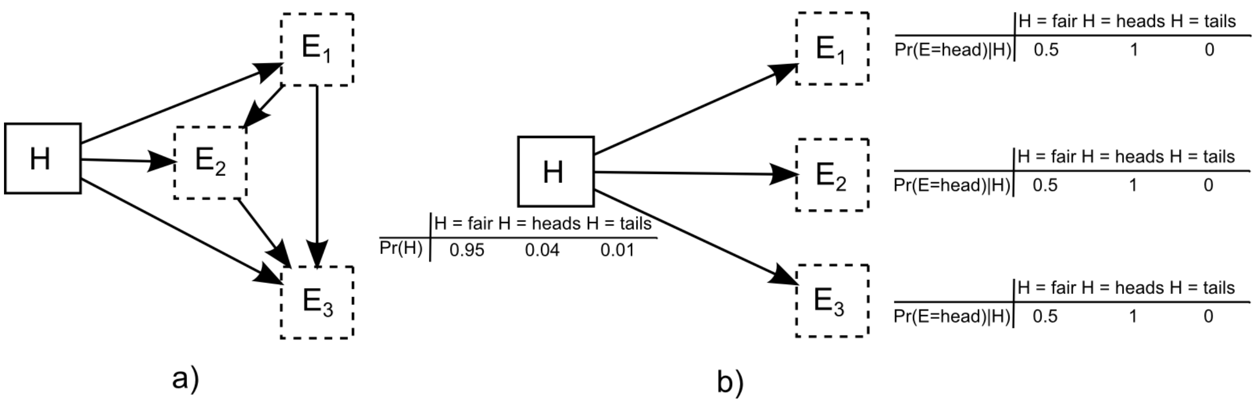

FIGURE 1: *Bayesian network with all dependencies (Fig. 1a) and only the direct dependencies (Fig. 1b) represented by directed arcs*

A Bayesian network is a correct representation of the problem domain if all variables $A_1, \ldots, A_{n-1}$ that have a *direct* influence on A_n are parents of A_n. If this is the case and the state of the parents of A_n is known, the variable A_n is independent of all its other ancestors. This means that in a Bayesian network,

$$\Pr(A_n|A_1, \ldots, A_{n-1}) = \Pr\big(A_n\big|\text{parents}(A_n)\big) \quad (1)$$

holds (for a proof, see (Charniak 1991)). This allows simplifying the chain rule, which in turn allows the calculation of the full joint probability distribution for the problem domain, to the following:

$$\Pr(A_1, \ldots, A_n) = \prod_{i=1}^{n} \Pr(A_i \mid A_1, \ldots, A_{i-1}) = \prod_{i=1}^{n} \Pr(A_i \mid \text{parents}(A_i)) \quad (7)$$

In the coin-tossing example, the full joint probability distribution can therefore be represented by

$$\mathbf{Pr}(H, E_1, E_2, E_3) = \Pr(H)\Pr(E_1|H)\Pr(E_2|H)\ \Pr(E_3|H)$$

which can be used to reconstruct the full joint probability distribution. If the subject in the coin-tossing example wishes to condition her belief in H on the evidence E_1, E_2, E_3, Bayes' rule tells her to calculate

$$\Pr(H|E_1, E_2, E_3) = \frac{\Pr(H, E_1, E_2, E_3)}{\Pr(\,E_1, E_2, E_3)} = \frac{\Pr(H)\Pr(E_1|H)\Pr(E_2|H)\ \Pr(E_3|H)}{\sum_H(H, E_1, E_2, E_3)}.$$

For demonstration purposes, the actual calculation is carried out for a very simple example, i.e., for the case where the subject observes that all three tosses of the coin land on heads. She must update her prior belief in the fairness of the coin, $\Pr(h_{fair}) = 0.95$, $\Pr(h_{heads}) = 0.04$, $\Pr(h_{tails}) = 0.01$, the following way (h = heads)

$$\Pr(h_{fair}|e_{1\,h}, e_{2\,h}, e_{3\,h}) = \frac{\Pr(h_{fair}, e_{1\,h}, e_{2\,h}, e_{3\,h})}{\Pr(e_{1\,h}, e_{1\,h}, e_{3\,h})}$$

$$= \frac{\Pr(h_{fair})\ \Pr(e_{1\,h}|h_{fair})\Pr(e_{2\,h}|h_{fair})\Pr(e_{3\,h}|h_{fair})}{\sum_H(e_{1\,h}, e_{1\,h}, e_{3\,h})}$$

$$= \frac{0.95 \times 0.5 \times 0.5 \times 0.5}{(0.95 \times 0.5 \times 0.5 \times 0.5) + (0.04 \times 1) + (0.01 \times 0)}$$

$$= 0.748.$$

That is, after observing three tosses that fall on heads in a row, her belief that the coin is fair is reduced from 0.95 to 0.75. For more complex queries, the calculation is tedious using paper and pencil even for small networks and impossible for large networks. Algorithms have been developed that perform these calculations efficiently for large networks (Pearl 1986; Lauritzen und Spiegelhalter 1988). For the user of Bayesian networks, knowledge of the algorithms is just as unnecessary as knowledge of the internal workings of a calculator is unnecessary for the use of a calculator (Fenton und Neil 2011). It is sufficient to know that the algorithms have been accepted by the scientific community as correct, and that different implementations lead to the same results. There are a number of both commercial and free software programs available for probabilistic inference using Bayesian networks. All calculations for this article were performed with SamIam (Sensitivity Analysis, Modeling, Inference And More) 3.0, which is developed by the Automated Reasoning Group of Professor Adnan Darwiche at UCLA. This software is free and well documented by a number of scientific papers and a book (Darwiche 2009); however, the same results could have been obtained by any number of programs.

C. The (lack of a) case for using Bayesian networks as decision aids by legal fact finders

While the introduction of Bayesian networks silences one concern of critics of the application of subjective probability theory to legal fact-finding, namely that it is not suitable for complex, real-world evaluations of evidence, it does not address the other concern, i.e. that the use of mathematical formal models leads to "wholly inaccurate, and misleadingly precise, conclusions" (Tribe 1971). Arguably, Bayesian decision networks even worsen the latter problem, because they require even more estimated probabilities that are not based on known relative frequencies, and they may convey an unjustified air of scientific accuracy. As explained in the introduction, the best case that can be made for any method of evidence evaluation is that it leads to an assessment of the case that more closely corresponds to the ground truth; however, this case can rarely be made for the simple reason that the ground truth is most often not known.

The second-best case for a method of evidence evaluation that can be made is that it avoids known fallacies of reasoning and common biases. Firstly because an irrational evaluation of evidence, i.e. an evaluation that violates rules of logic, is – at least in civil law countries – a violation of the rules of evidence *per se,* because evidence evaluation must be rational (Schweizer 2015). Secondly because it is reasonable to assume that a mental representation of the facts that is undistorted by logical fallacies and biases is more likely to correspond to reality. Proponents of the use of Bayesian networks for evidence evaluation in the legal context have argued that Bayesian networks may help avoid common fallacies of statistical reasoning (Fenton und Neil 2011; Berger 2014), but they have failed to submit the empirical evidence that would support their claim. Nobody seems to have claimed so far that the use of Bayesian networks may be beneficial beyond the avoidance of the well-known difficulties most people have with the intuitive application of Bayes' rule. If it could be shown that Bayesian networks help avoiding other common biases of reasoning, such as role-induced biases, the case for their use by legal fact-finders is strengthened.

IV. Experiment

A. Hypothesis

I hypothesize that forcing subjects to assess the likelihoods for each individual item of evidence and integrating the obtained values using a Bayesian network will lead to a reduction of role-induced bias in the assessment of the probability of guilt of a defendant in a criminal case. This is based on the observation that counterfactual thinking helps reduce coherence shifts (Simon 2004; O'Brien 2009). Subjects who consider why their hypothesis might be wrong show less bias (O'Brien 2009). When estimating the likelihood ratio that determines the probative value of a piece of evidence in a Bayesian framework, subjects are forced to estimate the probability that the evidence would be present if the hypothesis was wrong. This should induce counterfactual thinking that leads to a reduction of role-induced bias.

B. Method

a. Participants

I invited 89 subjects to the decision lab at the Max Planck Institute for Collective Goods, Bonn, Germany, for completion of a computer-based questionnaire. Sessions lasted about one hour. Subjects were paid € 10 for participation. The subjects were predominantly female (65%). Average age was 23 years (SD = 3.55). 14 (16%) studied law, with the rest divided among a number of fields including medicine, psychology, economics, sociology and agriculture.

b. Material and Procedure

The methodology of this study closely follows the one used by Schweizer 2014. The same instructions, case scenario and structure of the Bayes' net are employed. Key differences are the random assignment of the subjects to two groups intended to create a role-induced bias and the elicitation of not one probability estimate per parameter, but rather a lowest, highest and best estimate.

Subjects first read the instructions for completing the questionnaire and answered corresponding test questions before reading the case material. They were instructed that they could imagine a subjective probability of x% as the expectation of blindly drawing a red ball from an urn containing 100 balls, thereof x red balls. Additionally, the meaning of conditional probability was explained, and some examples were given that were not from the domain of legal evidence. It was explained that likelihoods need not add up to one. Subjects were asked to state probabilities as percentages from 0% to 100%, this being more natural than the mathematical convention of bounding probabilities between 0 and 1. Each explanation was followed by a test question. These were generally answered correctly by most participants (76% to 100% correct responses depending on the question). Subjects who gave a wrong answer were explained why they made a mistake.

Subjects were then randomly assigned to one of two conditions "defender" and "prosecutor", in which they were either told they were the accused's defense attorney or the state attorney prosecuting the case. They read an identical scenario of a case involving the (alleged) theft of money from a safe (the "Jason Wells/Hans H. case"), a scenario that has been used in a number of psychological studies (Simon et al. 2004; Engel und Glöckner 2013; Glöckner und Engel 2013; Schweizer 2014). The scenario was also handed out as a hard-copy and the participants were instructed that they could refer to the scenario while answering the questions.

The scenario, of just over 700 words, describes Hans, a 34-year-old married man with two children, employed at a construction company. Hans has recently been denied a promotion. He has a prior criminal record for attempted burglary at age 18 but has not since come into conflict with the law. One day, € 5,200 is missing from the company's safe. Eight people, among them Hans, have access to the safe, which was last opened at 7.14 pm. A technician testifies that he saw Hans leave the office in which the safe is located at about 7.15 pm. A surveillance video shows a car of the rare kind Hans drives leaving from the office building at 7.17 pm, but the license plate is illegible. Another witness, Silvia, testifies that she saw Hans at a school function at 8 pm wearing different clothes than the ones he wore at work, and it would be difficult to get from the office to the school in less than 40 min at that time of day. The day after the disappearance of the money from the safe, Hans repays a bank loan of € 4,870. Hans claims he received the money for the repayment from his sister-in-law who owns a flower shop, but he cannot produce a receipt for the transaction. He explains this by the practice in the flower business of "occasionally" doing business without issuing receipts.

After reading the scenario, subjects were asked to spend 10-15 minutes writing a plea to be presented to the court either arguing for the accused's guilt ("prosecutor" condition) or his innocence ("defender" condition). They then were asked to state the lowest, highest and best estimate for the subjective probability of the accused's guilt considering all the evidence. This procedure, known as "dialectical bootstrapping", is known to increase the accuracy of numerical estimates (Herzog und Hertwig 2009). In the context of estimating parameters for the computation of a Bayes' net, it has the additional advantage of avoiding – for the relevant best estimate – estimations of 0, which make the network incomputable and lead to a loss of data.

After reading the following definition of the criminal standard of proof used by the German Federal Supreme Court (e.g. BGH, 30 July 2009 – 3 StR 273/09 = BeckRS 2009, 25658; translation into English by the author)

> *"The conviction of the judge does not require an absolute certainty that excludes other possibilities with logical necessity. An adequate degree of certainty that overcomes reasonable doubt is sufficient. The judge is not prohibited from drawing possible, albeit not cogent inferences, from facts if such inferences are supported."*

they indicated the subjective probability they think this standard requires ("normative probability"), before stating whether as a judge, they would convict or acquit based on the evidence and the just read standard of proof.

They are then asked to indicate the lowest, highest and best estimate for the prior probability of guilt given that Hans is one of eight people with access to the safe, has been denied a promotion and has a prior criminal record. Subjects also indicated their prior belief in Hans having received the money from his sister-in-law (the other root node of the network), again indicating lowest, highest and best estimate.

The likelihood ratio for each item of evidence was elicited from the subjects using natural language questions. For example, for the witness statement of the technician, subjects must answer the questions

- "How likely is it that the technician testifies he saw Hans leaving the office, given that Hans left the office?"

and

- "How likely is it that the technician testifies he saw Hans leaving the office, given that Hans *did not* leave the office?",

in each case followed by "maximum: __%", "minimum: __%" and "best estimate: __%". Subjects assessed a total of nine likelihood ratios, which means they had to answer 18 questions, each requiring three probability estimates. At the end of the questionnaire, subjects were asked again what their subjective probability for Hans' guilt was and had to indicate which role they had been assigned to, to verify whether they correctly remembered their assigned role.

c. Computation of the Posterior Probability of Guilt

The posterior probability of guilt was computed for each subject using the "best estimate" for each prior and likelihood as obtained from that subject and the structure of the Bayesian network given in Figure 2. Evidence variables, i.e., variables the state of which is observed, are shown in dashed rectangles. The hypothesis variable is the variable of interest; it is shown in a rectangle with a thick border. Intermediate variables are variables that cannot be observed and mediate the influence of the evidence variables on the hypothesis variable; they are shown in solid line rectangles. The two evidence variables "refused promotion" and "old criminal record" are not required for the computation of the network, as their state is known and they are parents of the hypothesis variable. They are included in Figure 2 for sake of completeness.

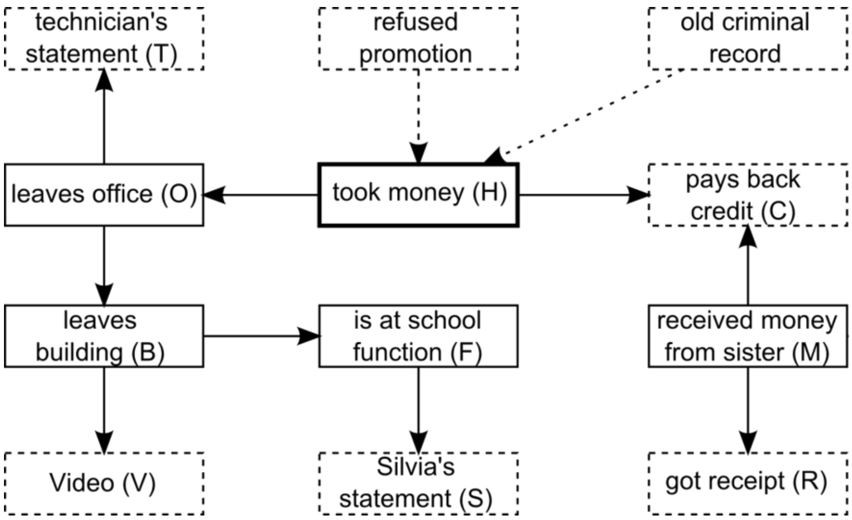

FIGURE 2: *Bayesian network for the Jason Wells/Hans H. case. Evidence variables with dashed borders, intermediate variables with solid borders, hypothesis variable with thick border*

Given the structure of the network in Figure 2, the full joint probability distribution for the network can be factorized as follows:

$$\Pr(H, M, O, B, F, T, V, S, R, C) = \Pr(H)\Pr(M)\Pr(O|H)\Pr(B|O)\Pr(F|B)\Pr(T|O)\Pr(V|B)\Pr(S|F)\Pr(R|M)\Pr(C|H, M)$$

The states of the evidence variables T, V, S and C are observed as true, and the state of the evidence variable R is observed as false. Known states are indicated with lower-case letters with the indices t = true and f = false. The query of interest is therefore

$$\Pr(h_t|M, O, B, F, t_t, v_t, s_t, r_f, w_f, c_t) = \frac{\Pr(h_t)\Pr(M)\Pr(O|h_t)\Pr(B|O)\Pr(F|B)\Pr(t_t|O)\Pr(v_t|B)\Pr(s_t|F)\Pr(r_f|M)\Pr(c_t|h_t, M)}{\sum_H \sum_M \sum_O \sum_B \sum_F \Pr(t_t, v_t, s_t, r_f, c_t)}$$

It would be highly impractical to carry out the actual summations required for the solution by hand, and all calculations were therefore performed using the software SamIam 3.0.

C. Results

Of the 89 subjects who participated in the experiment, one gave values for the likelihoods that made computation of the posterior for $\Pr(h_t)$ impossible (this participant answered all questions either with 0% or 100% and gave a holistic probability of 0% for the guilt of the accused. It is safe to say that he did not take the task seriously). He was excluded from further analysis as well as an additional eight subjects who did not remember correctly which role they had been assigned to when asked after they had answered all the questions, as it can be assumed that they cannot be properly influenced by role. This left a total of 80 subjects, 42 in the "defender" and 38 in the "prosecutor" condition.

46 subjects (57.5%) would have found the accused guilty of taking the money and 34 would have acquitted him if they had been the judge in the case. While the proportion of subjects in the "prosecutor" condition that would have found Hans guilty was slightly higher than that in the "defender" condition (60.5% versus 54.7%), the difference was not significant, χ^2 (1, N = 80) = 0.27, p = 0.6. The assigned role also did not lead to a significantly different expressed normative decision threshold required for a guilty verdict with M = 68.8, SE = 4.2 for "defenders" and M = 75.8, SE = 4.2 for "prosecutors", t(77.74) = – 1.17, p = 0.25.

Figure 3 shows probability of guilt after having read the scenario, but before answering any of the questions ("probability of guilt before"), the probability of guilt after answering the likelihood questions ("probability of guilt after") and the probability of guilt computed based on the answers to the likelihood questions using the Bayesian network exposed above ("computed probability of guilt") according to assigned role. The horizontal dashed line indicates the mean decision threshold ("standard of proof") for both treatment groups combined (M = 72.1, SE = 3.0).

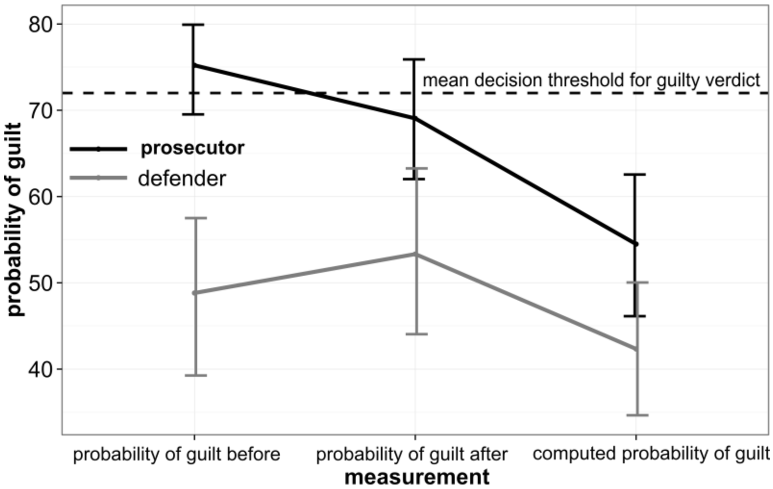

FIGURE 3: *measurements of probability of defendant's guilt before and after assessing the probative value of the evidence using likelihood-ratios, and probability of guilty as computed by Bayesian network*

Since the three measurements of probability of guilt are not independent, I fit a multilevel model using the R package "lmerTest" (Kuznetsova et al. 2016). In detail, I used two centered dummies for measurement with condition "probability of guilt before" as reference group, one centered dummy for role with condition "defender" as reference group and two interaction terms between each dummy for measurement and dummy for role. To account for dependencies between repeated measures, I included a random intercept depending on participants.

The main effect of being assigned to the role of prosecutor versus defendant on the probability of guilt is large and significant, $b = 19.13$, $t(78) = 4.05$, $p < .001$. Answering the likelihood questions has no significant main effect on probability of guilt in comparison to the probability of guilt before, $b = -0.55$, $t(156) = -0.18$, $p = .86$. Computing the probability of guilt using the Bayesian network parametrized with the answers from the likelihood questions leads to a significantly reduced probability of guilt in comparison to the probability of guilt before, $b = -12.95$, $t(156) = -4.19$, $p < .001$.

The interaction of role and probability of guilt after answering the likelihood questions is marginally significant, $b = -10.56$, $t(156) = -1.71$, $p = 0.089$, while the interaction of role and computed probability of guilt is significant, $b = -14.63$, $t(156) = -2.36$, $p < .05$. This means that the difference in ratings of guilt between prosecutor and defender is (marginally) significantly lower after participants evaluated the evidence and significantly lower when computed. Thus, prosecutors and defenders are more equal in probability ratings of guilt after rating the evidence or after computing. Nonetheless, ratings of probability of guilt between prosecutors and defenders are still significantly different from each other after rating the evidence ($t(79) = -2.78$, $p < .01$, two-tailed) and when computed ($t(78) = -2.05$, $p = .043$, two-tailed). Thus, although the role induced bias significantly reduces as shown by the interaction terms after rating the evidence and when computed, the role induced bias is still present as shown in the two post-hoc comparisons.

Table 3 in the appendix reports the likelihoods for each individual item of evidence and each intermediate hypothesis, according to group. As the computed posterior probability of guilt does not fully converge, it is of interest to examine whether the defenders and prosecutors differ in the assessment of specific items of evidence. As Table 3 two shows, there are only two likelihoods that are assessed significantly differently by the two groups: the prosecutors consider it more likely that the accused could be at the school function in time although he left the office building at 7:17 pm (M = 53.8, SE = 4.8, versus M = 39.1, SE = 5.1, $t(77.99) = -2.08$, $p = .041$). On average, prosecutors also consider it more likely that the accused would have a receipt if he had in fact received the money from his sister-in-law as claimed (M = 50.4, SE = 5.3, versus M = 35.3, SE = 4.3 for the defenders, $t(73.44) = -2.23$, $p = .03$).

In a Bayesian framework, the probative value of an item of evidence is determined not by the likelihoods per se, but by their ratio. Figure 4 therefore shows the likelihood *ratios* for each item of evidence and each intermediate hypothesis. The values in Figure 4 were calculated by dividing the mean values from Table 3 in the appendix for each group and likelihood. E.g., the value for the likelihood ratio for the accused leaving the office shortly after 7:15 pm for the defenders is calculated by dividing the mean likelihood that the accused leaves the office shortly after 7:15 pm, given that he took the money, and the mean likelihood that the accused leaves the office shortly after 7:15 pm, given that he did not take the money. The specific value for the defenders for this likelihood ratio is 59.9/46.8 = 1.28, and for the prosecutors 75.8/40.0 = 1.90. Likelihood-ratios above 1 mean that the evidence is incriminating for Hans. In other words, the prosecutors consider leaving the office shortly after 7:15 pm slightly more incriminating than the defenders. Likelihood-ratios below 1 mean that the evidence is exculpating for Hans. To make the graph in Figure 4 more legible, likelihood ratios below 1 were converted into negative values by taking the inverse and adding a negative sign (e.g., a likelihood ratio of 0.5 would be displayed as $-1/0.5 = -2$).

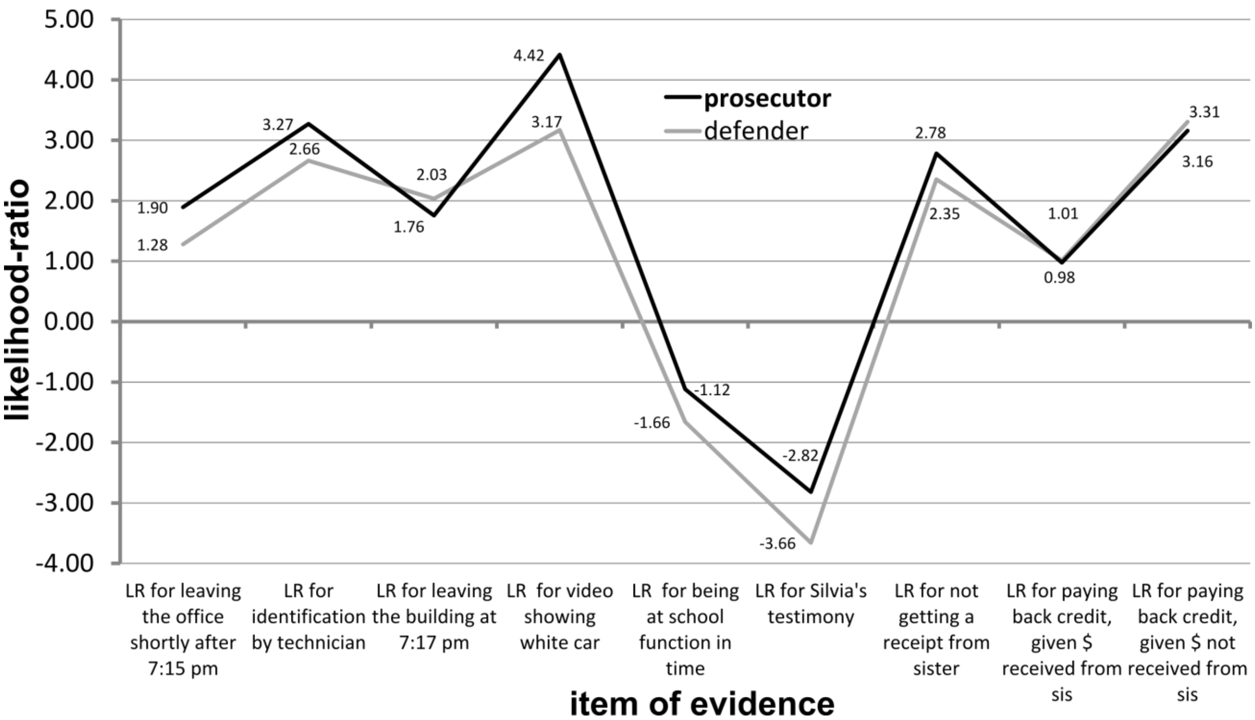

FIGURE 4: *likelihood-ratios for each individual piece of evidence, as assessed by prosecutors and defenders*

Figure 4 shows graphically that the he probative force of the evidence is assessed very similarly by both groups, i.e., both groups tend to agree with the direction and strength of the probative value. The participants assigned the role as prosecutors do, however, assess the incriminating evidence as stronger than the defenders, and the exculpating evidence as less probative. Therefore, the line for the average likelihood ratios of the prosecutors is, with the exception of the likelihood-ratio for leaving the building at 7:17 pm, always above the line for the defenders. Interestingly, both groups agree that the observation that the accused paid back a bank credit the day after the money disappeared from the safe is of no probative value if he received money from his sister (likelihood-ratio ≈ 1). The timing of the loan repayment is, however, incriminating if the accused did not receive money from his sister (average likelihood-ratio slightly above 3).

Participants should only convict if their posterior probability of guilt meets or exceeds the threshold value required for a conviction. Table 2 therefore compares the probabilities for guilt before, after and the computed probability of guilt (column headings) with the probability threshold required for a conviction (line headings) and counts instances where the posterior exceeds the threshold. Two different threshold values are used: In the first line the average interpretation, as a posterior subjective probability, of the Federal Court of Justice's verbal description of the standard of proof in criminal cases (which is M = 72.1, SE = 3.0). In the second line the probability of guilt is compared to each subject's own interpretation of the court's standard ("personal threshold exceeded").

TABLE 2: *Instances of posteriors meeting or exceeding the threshold probability for conviction*

	Probability for guilt (before)		Probability for guilt (after)		Computed probability for guilt	
	defenders (n=42)	prosecutors (n=38)	defenders (n=42)	prosecutors (n=38)	defenders (n=42)	prosecutors (n=38)
Avg. threshold exceeded	13 (31%)	22 (57%)	15 (36%)	20 (52%)	6 (14%)	10 (26%)
Personal threshold exceeded	9 (21%)	13 (34%)	12 (29%)	8 (21%)	8 (19%)	11 (29%)

If the posterior exceeds the personal threshold value, the subject should (according to his or her own standard) convict or, if not, acquit. The top left cell of Table 2 shows that the probability of guilt (before answering the likelihood questions) for only 13 subjects in the "defenders" condition exceeds the average required threshold value for a conviction, versus 22 subjects in the "prosecutors" condition, χ^2 (1, N = 80) = 5.21, p = .02. All the other differences between "defenders" and "prosecutors" in Table 2 are not statistically significant at a level of $p \leq .05$.

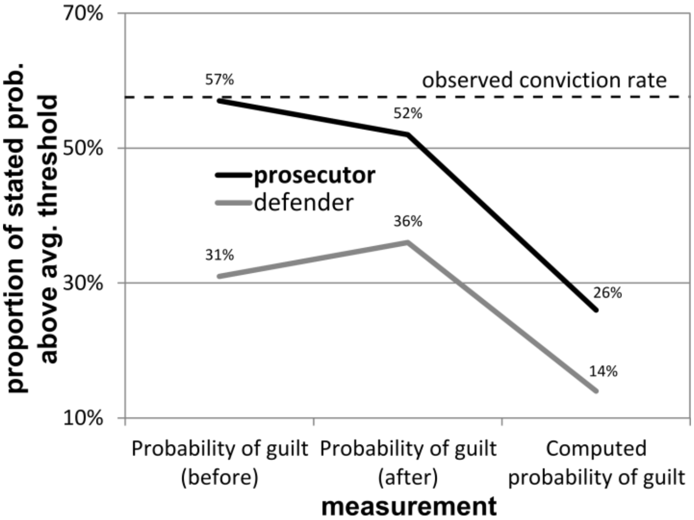

FIGURE 5: *proportions of probabilities of guilt exceeding the mean stated normative threshold for conviction*

Figure 5 is a visualization of the first line of Table 2, showing the proportions of stated and computed probabilities of guilt exceeding the average required decision threshold. 57.5% (46) of all the subjects would have found the accused guilty had they been the judge in the case. As is evident, only a fraction of those 46 subjects *should have* convicted had they adhered to their *own* understanding of the standard of proof. The dashed horizontal line in Figure 5 is drawn at 57.5% to indicate the proportion of observed convictions.

D. Discussion

The experimental data support the hypothesis that thinking in a likelihood-framework, which requires considering the probability that the evidence would exist assuming the accused's innocence, reduces the extent of role-induced bias. It is not sufficient to make it disappear, though. Computing the posterior probability of guilt according to the constraints imposed on the partial beliefs by the axioms of subjective probability theory leads to a significant and large further reduction of the role-induced bias.

One, admittedly speculative, reason why the computation of the network – parametrized with each subjects own partial beliefs – still further reduces role-induced bias is that the scenario contains (corroboratively) redundant items of evidence. "Redundant items of evidence" are two or more items of evidence that are interdependent in the sense that taking into account the first item of evidence lowers the probative value of the second item of evidence, it makes the second piece "redundant" (Lempert 1977) Lempert was the first to hypothesize that humans intuitively overvalued redundant evidence (Lempert 1977), and Schum and Martin provided empirical evidence for the proposition (Schum und Martin 1982). In the scenario here used, the testimony of the technician to the effect that he saw somebody he identified as the accused leaving the office shortly after the alleged theft and the video footage showing a car of the type, colour and model driven by the accused getting away from the office building immediately afterwards are redundant in the sense that they both only support the proposition that the accused was near the safe at the time of the alleged theft. If one accepts this proposition based on one of the pieces of evidence, the probative value of the additional item of evidence is nil. The evaluation of the evidence using the structure of the Bayesian network makes this immediately clear and avoids the mistake that may intuitively be made.

There are no significant differences in the expression of the normative standard of proof as a threshold probability between the defenders and the prosecutors. However, the subjects are inconsistent with their own standard: 24 subjects convicted the accused although their own stated probability of guilt posterior did not meet their own definition of the legal standard of proof expressed as a subjective probability (see Table 3).

A limitation of this study is that the structure of the network was designed by the experimenter and therefore the same for all subjects. When evaluating a mass of evidence using a Bayesian network, the impact of the evidence on the probability of the hypothesis of interest depends not only on the probative force of each individual item of evidence, but also on their interrelations. In other words, the assessment of the overall probative force of the evidence depends (also) on the structure of the network. As the direct dependencies which structure the model are based on the expert's knowledge and assumptions about the workings of the world, different experts may structure the problem differently. Further research should explore whether the further large reduction in the difference of the probabilities of guilt between the defenders and prosecutors are still to be found if not only the parameters for the network, but also the network *structure* is elicited from the subjects.

V. Conclusion

This study empirically demonstrates an advantage of using a Bayesian network for the evaluation of evidence in a case where there are no relative frequencies that could form the basis for assessing the probative value of the evidence. The study shows that the large effect role has on the assessment of the accused's guilt is reduced when the evidence is assessed in a likelihood-framework and the partial beliefs are integrated according to the axioms of probability theory. Since role-induced bias may lead to unnecessary indictments and unnecessary expenditure of public funds, the study supports the case for using Bayesian networks as decision aids for evidence evaluation in the legal context.

Appendix

TABLE 2: *Mean likelihoods, by group (standard deviation)*

	$Pr(o_t \mid h_t)$	$Pr(o_t \mid h_f)$	$Pr(t_t \mid o_t)$	$Pr(t_t \mid o_f)$	$Pr(b_t \mid o_t)$	$Pr(b_t \mid o_f)$	$Pr(v_t \mid b_t)$	$Pr(v_t \mid b_f)$
Defender	59.9 (29.7)	46.8 (28.7)	83.9 (19.8)	31.5 (25.5)	72.2 (24.7)	35.5 (22.2)	84.0 (28.8)	26.5 (29.7)
Prosecutor	75.8 (26.3)	40.0 (28.8)	91.3 (24.6)	27.9 (23.9)	75.2 (21.3)	42.8 (28.3)	87.0 (21.9)	19.7 (29.4)
Difference	15.9*	6.8	7.4	3.6	2.6	7.3	3.0	6.8

** $p<0.05$, * $p<0.1$ (using a two sided t-test).

TABLE 2 (CONT'D): *Mean likelihoods, by group (standard deviation)*

	$Pr(f_t \mid b_t)$	$Pr(f_t \mid b_f)$	$Pr(s_t \mid f_t)$	$Pr(s_t \mid f_f)$	$Pr(r_t \mid m_t)$	$Pr(r_t \mid m_f)$
Defender	39.1 (33.4)	64.8 (29.0)	86.7 (19.9)	23.7 (21.9)	35.3 (27.9)	15.0 (21.6)
Prosecutor	53.8 (29.9)	60.1 (31.4)	89.3 (26.7)	31.7 (27.0)	50.4 (32.4)	18.1 (25.7)
Difference	14.7**	4.7	2.6	8.0	15.1**	3.1

TABLE 2 (CONT'D): *Mean likelihoods, by group (standard deviation)*

	$Pr(c_t \mid h_t, m_t)$	$Pr(c_t \mid h_f, m_t)$	$Pr(c_t \mid h_t, m_f)$	$Pr(c_t \mid h_f, m_f)$
Defender	47.6 (35.4)	47.1 (26.2)	63.8 (28.5)	19.3 (25.1)
Prosecutor	54.8 (33.8)	56.1 (30.5)	63.8 (27.9)	20.2 (25.9)
Difference	7.2	9.0	0.0	0.9

References

Aitken, Colin G. (2012): An introduction to a debate. In: *Law, Probability & Risk* (11), S. 255–258. DOI: 10.1093/lpr/mgs014.

Allen, Ronald J. (1991): The Nature of Juridical Proof. In: *Cardozo Law Review* 13, S. 373–444.

Amaya, Amalia (2008): Justification, Coherence, and Epistemic Responsibility in Legal Fact-Finding. In: *Episteme* 5 (3), S. 306–319. DOI: 10.3366/E1742360008000415.

Anderson, Craig A.; Lepper, Mark R.; Ross, Lee (1980): Perseverance of social theories. The role of explanation in the persistence of discredited information. In: *Journal of Personality and Social Psychology* 39 (6), S. 1037.

Berger, Daniel (2014): Improving Legal Reasoning using Bayesian Probability Methods. PhD Thesis. London: Queen Mary University of London.

Brownstein, Aaron L. (2003): Biased predecision processing. In: *Psych. Bull.* 129 (4), S. 545–568. DOI: 10.1037/0033-2909.129.4.545.

Byrne, Michael D. (1995): The Convergence of Explanatory Coherence and the Story Model. A Case Study in Juror Decision. In: Johanna D. Moore und Jill Fain Lehman (Hg.): Proceedings Of The Seventeenth Annual Conference Of The Cognitive Science Society. Mahwah, N.J: Lawrence Erlbaum, S. 539–549.

Callen, Craig R. (1982): Notes on a grand illusion: some limits on the use of Bayesian theory in evidence law. In: *Indiana Law Journal* 57, S. 1–44.

Carlson, Kurt A.; Russo, J. Edward (2001): Biased interpretation of evidence by mock jurors. In: *Journal of Experimental Psychology: Applied* 7 (2), S. 91–103. DOI: 10.1037/1076-898X.7.2.91.

Charniak, Eugene (1991): Bayesian Networks without Tears. In: *AI Magazine* 12 (4), S. 50–61.

Christensen, David (2007): Putting logic in its place. Formal constraints on rational belief. Oxford: Clarendon Press. Online verfügbar unter http://www.gbv.de/dms/bowker/toc/9780199204311.pdf.

Cohen, L. Jonathan (1977): The probable and the provable. Oxford: Clarendon Press.

Darwiche, Adnan (2009): Modeling and reasoning with Bayesian networks. Cambridge, New York: Cambridge Univ. Press.

de Finetti, Bruno (1937): La prévision: ses lois logiques, ses sources subjectives. In: *Annales de l'Institut Henri Poincaré* 7 (1), S. 1–68.

DeKay, Michael L. (2015): Predecisional Information Distortion and the Self-Fulfilling Prophecy of Early Preferences in Choice. In: *Current Directions in Psychol Sci* 24 (5), S. 405–411. DOI: 10.1177/0963721415587876.

Edwards, Ward (1991): Influence Diagrams, Bayesian Imperialism, and the Collins Case: An Appeal to Reason. In: *Cardozo Law Review* 13, S. 1025–1074.

Engel, Christoph; Glöckner, Andreas (2013): Role-Induced Bias in Court. An Experimental Analysis. In: *Journal of Behavioral Decision Making* 26 (3), S. 272–284. DOI: 10.1002/bdm.1761.

Fenton, Norman; Neil, Martin (2011): Avoiding Probabilistic Reasoning Fallacies in Legal Practice using Bayesian Networks. In: *Australian Journal of Legal Philosohpy* 36, S. 114–151. Online verfügbar unter http://heinonline.org/HOL/Page?handle=hein.journals/ajlph36&id=118&collection=journals&index=journals/ajlph.

Fenton, Norman; Neil, Martin; Berger, Daniel (2016): Bayes and the Law. In: *Annual review of statistics and its application* 3, S. 51–77. DOI: 10.1146/annurev-statistics-041715-033428.

Fenton, Norman; Neil, Martin; Lagnado, David A. (2013): A General Structure for Legal Arguments About Evidence Using Bayesian Networks. In: *Cognitive Science* 37 (1), S. 61–102. DOI: 10.1111/cogs.12004.

Glöckner, Andreas; Betsch, Tilmann (2008): Modeling option and strategy choices with connectionist networks. Towards and integrative model of automatic and deliberate decision making. In: *Judgment and Decision Making* 3 (3), S. 215–228.

Glöckner, Andreas; Engel, Christoph (2013): Can We Trust Intuitive Jurors? Standards of Proof and the Probative Value of Evidence in Coherence-Based Reasoning. In: *Journal of Empirical Legal Studies* 10 (2), S. 230–252. DOI: 10.1111/jels.12009.

Good, Irving John (1950): Probability and the weighing of evidence. London: Charles Griffin.

Goodman-Delahunty, Jane; Granhag, Pär Anders; Hartwig, Maria; Loftus, Elizabeth F. (2010): Insightful or wishful: Lawyers' ability to predict case outcomes. In: *Psychology, Public Policy, and Law* 16 (2), S. 133–157. DOI: 10.1037/a0019060.

Herzog, Stefan M.; Hertwig, Ralph (2009): The wisdom of many in one mind improving individual judgments with dialectical bootstrapping. In: *Psychological Science* 20 (2), S. 231–237.

Holyoak, Keith J.; Simon, Dan (1999): Bidirectional reasoning in decision making by constraint satisfaction. In: *Journal of Experimental Psychology: General* 128 (1), S. 3–31. DOI: 10.1037/0096-3445.128.1.3.

Hope, Lorraine; Memon, Amina; McGeorge, Peter (2004): Understanding Pretrial Publicity: Predecisional Distortion of Evidence by Mock Jurors. In: *Journal of Experimental Psychology: Applied* 10 (2), S. 111–119. DOI: 10.1037/1076-898X.10.2.111.

Janis, Irving L.; King, Bert (1954): The influence of role playing on opinion change. In: *Journal of abnormal psychology* 49 (2), S. 211–218.

Juchli, Patrick; Biedermann, Alex; Taroni, Franco (2012): Graphical probabilistic analysis of the combination of items of evidence. In: *Law, Probability & Risk* 11 (1), S. 51–84. DOI: 10.1093/lpr/mgr023.

Kadane, Joseph B.; Schum, David A. (1996): A probabilistic analysis of the Sacco and Vanzetti evidence. New York: Wiley. Online verfügbar unter http://catdir.loc.gov/catdir/bios/wiley041/95042458.html.

Kahneman, Daniel; Tversky, Amos (1973): On the psychology of prediction. In: *Psych. Rev.* 80 (4), S. 237–251. DOI: 10.1037/h0034747.

Kaye, David H. (1979): The Laws of Probability and the Law of the Land. In: *University of Chicago Law Review* 47 (1), S. 34–56. Online verfügbar unter http://www.jstor.org/stable/1599414.

Kunda, Ziva; Thagard, Paul (1996): Forming impressions from stereotypes, traits, and behaviors: A parallel-constraint-satisfaction theory. In: *Psych. Rev.* 103 (2), S. 284–308. DOI: 10.1037/0033-295X.103.2.284.

Kuznetsova, Alexandra; Brockhoff, Per Bruun; Christensen, Rune Haubo (2016): lmerTest: Tests in Linear Mixed Effects Models. R package version 2.0-32. Online verfügbar unter https://CRAN.R-project.org/package=lmerTest.

Lagnado, David A.; Fenton, Norman; Neil, Martin (2012): Legal idioms: a framework for evidential reasoning. In: *Argument & Computation* 4 (1), S. 46–63. DOI: 10.1080/19462166.2012.682656.

Lauritzen, Steffen L.; Spiegelhalter, David J. (1988): Local Computations with Probabilities on Graphical Structures and Their Application to Expert Systems. In: *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* 50 (2), S. 157–224. Online verfügbar unter http://www.jstor.org/stable/2345762.

Lempert, Richard O. (1977): Modeling Relevance. In: *Michigan Law Review* 75 (5/6), S. 1021–1057. Online verfügbar unter http://www.jstor.org/stable/1288024.

Lundberg, C. Gustav (2004): Modeling and predicting emerging inference-based decisions in complex and ambiguous legal settings. In: *European Journal of Operational Research* 153 (2), S. 417–432. DOI: 10.1016/S0377-2217(03)00162-0.

Monroe, Brian M.; Read, Stephen J. (2008): A general connectionist model of attitude structure and change: The ACS (Attitudes as Constraint Satisfaction) model. In: *Psychological Review* 115 (3), S. 733.

O'Brien, Barbara (2009): Prime suspect: An examination of factors that aggravate and counteract confirmation bias in criminal investigations. In: *Psychology, Public Policy, and Law* 15 (4), S. 315.

Park, Roger C.; Tillers, Peter; Moss, Frederick C.; Risinger, D. Michael; Kaye, David H.; Allen, Ronald J. et al. (2010): Bayes Wars Redivivus. An Exchange. In: *International Commentary on Evidence* 8 (1), S. 1–38. DOI: 10.2202/1554-4567.1115.

Pearl, Judea (1986): Fusion, Propagation, and Structuring in Belief Networks. In: *Artificial Intelligence* 29, S. 241–288.

Pearl, Judea (1988): Probabilistic reasoning in intelligent systems. Networks of plausible inference. San Francisco, Calif: Morgan Kaufmann.

Pennington, Nancy; Hastie, Reid (1988): Explanation-based decision making. Effects of memory structure on judgment. In: *Journal of Experimental Psychology: Learning, Memory, and Cognition* 14 (3), S. 521–533. DOI: 10.1037/0278-7393.14.3.521.

Read, Stephen J.; Vanman, Eric J.; Miller, Lynn C. (1997): Connectionism, Parallel Constraint Satisfaction Processes, and Gestalt Principles. (Re) Introducing Cognitive Dynamics to Social Psychology. In: *Personality and Social Psychology Review* 1 (1), S. 26–53. DOI: 10.1207/s15327957pspr0101_3.

Robertson, Bernard; Vignaux, G. A. (1992): Taking fact analysis seriously. In: *Michigan Law Review* 91, S. 1442–1464.

Schum, David A. (1988): Probability and the processes of discovery, proof, and choice. In: Peter Tillers (Hg.): Probability and inference in the law of evidence. The uses and limits of Bayesianism. Dordrecht: Reidel (Boston studies in the philosophy of science, 109), S. 213–270.

Schum, David A.; Martin, Anne W. (1982): Formal and Empirical Research on Cascaded Inference in Jurisprudence. In: *Law & Society Review* 17 (1), S. 105–151. Online verfügbar unter http://www.jstor.org/stable/3053534.

Schweizer, Mark (2014): Comparing holistic and atomistic evaluation of evidence. In: *Law, Probability & Risk* 13 (1), S. 65–89. DOI: 10.1093/lpr/mgt013.

Schweizer, Mark (2015): Beweiswürdigung und Beweismaß. Rationalität und Intuition. Tübingen: Mohr Siebeck (Jus Privatum, 189).

Simon, Dan (2004): A Third View of the Black Box. Cognitive Coherence in Legal Decision Making. In: *University of Chicago Law Review* 71 (2), S. 511–586. Online verfügbar unter http://www.jstor.org/stable/1600674.

Simon, Dan; Krawczyk, Daniel C.; Holyoak, Keith J. (2004): Construction of Preferences by Constraint Satisfaction. In: *Psych. Science* 15 (5), S. 331–336. DOI: 10.1111/j.0956-7976.2004.00678.x.

Taroni, Franco; Aitken, Colin G.; Garbolino, Paolo; Biedermann, Alex (2006): Bayesian networks and probabilistic inference in forensic science. Chichester: Wiley (Wiley series on statistics in practice). Online verfügbar unter http://www.loc.gov/catdir/enhancements/fy0658/2005057711-d.html / http://www.gbv.de/dms/hbz/toc/ht014696644.pdf.

Thagard, Paul (1989): Explanatory coherence. In: *Behavioral and Brain Sciences* 12, S. 435–467.

Thagard, Paul (2003): Why wasn't O.J. convicted? Emotional coherence in legal inference. In: *Cognition & Emotion* 17 (3), S. 361–383. DOI: 10.1080/02699930244000002.

Thagard, Paul; Verbeurgt, Karsten (1998): Coherence as constraint satisfaction. In: *Cognitive Science* 22 (1), S. 1–24. DOI: 10.1016/S0364-0213(99)80033-0.

Thompson, Leigh; Loewenstein, George (1992): Egocentric interpretations of fairness and interpersonal conflict. In: *Organizational Behavior and Human Decision Processes* 51 (2), S. 176–197.

Tillers, Peter (Hg.) (1988): Probability and inference in the law of evidence. The uses and limits of Bayesianism. Symposium on Probability and Inference in the Law of Evidence. Dordrecht: Reidel (Boston studies in the philosophy of science, 109).

Tribe, Laurence H. (1971): Trial by Mathematics. Precision and Ritual in the Legal Process. In: *Harvard Law Review* 84 (6), S. 1329–1393. Online verfügbar unter http://www.jstor.org/stable/1339610.

Tversky, Amos; Kahneman, Daniel (1982): Evidential impact of base rates. In: Daniel Kahneman, Paul Slovic und Amos Tversky (Hg.): Judgment under uncertainty. Heuristics and biases. Cambridge,, New York: Cambridge Univ. Press, S. 153–160.

Tversky, Amos; Kahneman, Daniel (1983): Extensional versus intuitive reasoning. The conjunction fallacy in probability judgment. In: *Psych. Rev.* 90 (4), S. 293–315. DOI: 10.1037/0033-295X.90.4.293.

Vidmar, Neil; Laird, Nancy M. (1983): Adversary social roles. Their effects on witnesses' communication of evidence and the assessments of adjudicators. In: *Journal of Personality and Social Psychology* 44 (5), S. 888–898. DOI: 10.1037/0022-3514.44.5.888.

Zimbardo, Philip G. (1965): The effect of effort and improvisation on self-persuasion produced by role-playing. In: *Journal of Experimental Social Psychology* 1 (2), S. 103–120. DOI: 10.1016/0022-1031(65)90039-9.

Zoete, Jacob de; Sjerps, Marjan; Lagnado, David; Fenton, Norman (2014): Modelling crime linkage with Bayesian networks. In: *Science & Justice* 46, S. 209–217. DOI: 10.1016/j.scijus.2014.11.005.