university of groningen / faculty of science and engineering / artificial intelligence

Modelling Criminal Fraud Cases in Bayesian Networks

Yasmin de Groot
Supervisors: L. S. van Leeuwen & prof. dr. H.B. Verheij

Abstract: This thesis examines the application of Bayesian networks in modelling criminal cases, evaluating the benefits and limitations of the scenario idiom in a new legal case. While earlier research has successfully applied these methods to murder cases, this study assesses their effectiveness in fraud cases. This research uses scenario-based Bayesian networks to analyse the Dotterbloem case, involving a former Ministry of Defence employee convicted of passive corruption and breach of secrecy. The findings highlight several benefits, including how scenario schemes provide guidance in the interactions of elements, as well as the ability to accommodate various forms of evidence. However, the study also identifies limitations, including challenges in merging different narratives and determining accurate prior probabilities.

1 Introduction

Ensuring the reliability and transparency of evidence interpretation in legal judgments is crucial to maintaining the integrity of the judicial process (Hans & Saks, 2018). According to a study by Desai et al. (2016), the application of Bayesian networks in legal judgments improves both the reliability and transparency of evidence interpretation. This thesis focusses on Bayesian networks in conjunction with narrative idioms. These narrative idioms, a structured approach to evidence interpretation, have previously been applied successfully in the context of murder cases, as demonstrated by Vlek et al. (2016). This research aims to use these combined methods to investigate fraud cases.

1.1 Previous research

To formalise the theory of narrative coherence in legal contexts, Vlek et al. (2016) introduced the narrative idioms. This methodology aims to capture the coherence of a scenario within a Bayesian network. The narrative idioms are based on the premise that the elements of a scenario collectively form a coherent whole. Vlek et al. (2014) conducted a case study in which Bayesian networks incorporating narrative idioms were successfully applied to murder cases.

The influence of storytelling is of great significance for legal judgements, according to Pennington and Hastie 1991. In their research, it was shown that jurors create stories based on evidence, helping them understand and remember the facts. The created narrative framework helps jurors organise diverse pieces of evidence into a coherent whole. Jurors evaluate the completeness, consistency, and coherence of their stories to make a judgment.

The approach developed by Vlek and other previous studies (Pennington & Hastie (1991); Fenton et al. (2013)) has primarily focused on modelling murders. The capability to evaluate forensic evidence using statistical probabilities aligns well with Bayesian networks, making it straightforward to represent entire cases within Bayesian networks. However, this prompts significant questions about the adaptability of this method to other types of crime, such as fraud.

Fraud cases often involve ambiguous actions and lack physical evidence, relying heavily on the interpretation of intent rather than tangible proof. This is particularly challenging because intent must be inferred from circumstantial evidence, such as statements, conduct, and witness testimony. According to Kammen & Moudy (2023), the proof of intent in fraud cases does not require direct evidence, but rather the accumulation of circumstantial evidence that collectively indicates intent beyond reasonable doubt.

Unlike murder, which typically has clear, direct actions leading to an outcome, fraud can be passive or active, involving complex layers of deception and often lacking physical evidence. These distinctions raise questions about whether Bayesian networks using the scenario idiom need adjustments or a reconfigured approach to accurately accommodate and analyse the subtleties and complexities of fraudulent activities.

This thesis addresses the specific challenges associated with fraud cases and the interpretation of complex evidence in fraud. The purpose of the study is to contribute to a better understanding of how this scenario idiom method can be integrated into legal proceedings to support clearer and more reliable outcomes.

1.2 Research question

The research question for this thesis will be: *To what extent could a scenario-based Bayesian network method be applied to fraud analysis?*

A case will be modelled in such a scenario-based Bayesian network. In order to do this, scenarios are created from the Dotterbloem case, which will be modelled using the methods described in (Vlek et al., 2016). This Dotterbloem case concerns a former employee at the Ministry of Defence, convicted of passive corruption and breach of secrecy. More information on the case (in Dutch) can be found at (Rechtbank Rotterdam, 2018b).

The application of this method to fraud investigations is expected to reveal certain limitations, given the differences in the nature of evidence between fraud and murder cases.

2 Theoretical background

In this section, we introduce Bayesian networks, theory on the construction of stories, and how these two were combined to create the scenario idiom.

2.1 Bayesian networks

A Bayesian network is a compact representation of a joint probability distribution Pr over a set of variables $\mathbf{V}$. This network is structured as a directed acyclic graph (DAG), denoted as G, where each node represents a variable and is associated with a conditional probability table (CPT) (Jensen & Nielsen, 2007). The dependencies between variables are modelled by directed edges E. These edges may represent causal relationships, but are not necessarily causal (Dawid, 2010). Together, the variables and edges define the graph G and the Bayesian network: $\langle G, \Pr\rangle = \langle \mathbf{V}, E, \Pr\rangle$. From the structure of the network, it can be read which variables possibly have an influence on each other.

In a Bayesian network, each node X includes a CPT, which reveals the probability that X occurs based only on the possible combinations of values from its parent nodes in the network. Essentially, the CPT quantifies the dependency of X on its parent nodes. If X has parent nodes $P_a(X)$, the conditional probability of X is represented as $\Pr(X|P_a(X))$. Using these tables, any probability of interest can be calculated applying the chain rule for Bayesian networks, which states that the joint probability distribution over all variables $X_1, X_2, \ldots, X_n$ can be factored as:

$$\Pr(X_1, X_2, \ldots, X_n) = \prod_{i=1}^{n} \Pr(X_i|P_a(X_i)) \qquad (2.1)$$

One use case for Bayesian networks is their ability to compactly represent a joint probability distribution. Instead of enumerating all possible combinations of variables and their associated probabilities, which can be computationally intensive, the network focusses on the conditional probabilities of each node relative to its parent nodes. This compact representation simplifies the computation of probabilities and the updating of beliefs when new evidence is introduced. For example, when new evidence E is observed, the belief in any node X can be updated using Bayes' theorem:

$$\Pr(X|E) = \frac{\Pr(E|X)\Pr(X)}{\Pr(E)} \qquad (2.2)$$

This process, known as belief propagation or probabilistic inference, allows Bayesian networks to efficiently handle complex probabilistic reasoning tasks (Pearl, 1988).

2.2 Structures of stories

The way a story is structured greatly impacts its perceived plausibility (Pennington & Hastie, 1991).

This is particularly clear in legal contexts, where prosecutors often try to create a coherent sequence of events that would have caused the evidence provided. The human cognitive tendency to organise information into narrative structures means that when the events of a case are presented in chronological order and contain a central action, they are perceived as more credible and convincing.

In their work, Pennington and Hastie 1991 developed the *story model* of juror decision making, which suggests that jurors create stories to make sense of the evidence and testimonies they encounter in court. According to this model, jurors do not assess individual pieces of evidence in isolation. Instead, they integrate these pieces into a comprehensive narrative that explains the sequence of events leading to the suggested outcome. This narrative construction allows jurors to fill in gaps, make inferences, and establish causal links between events, thus improving their understanding and retention of the case details.

The principles that determine a story's acceptability and confidence level are known as certainty principles (Pennington & Hastie, 1991). According to the theory, two certainty principles influence acceptance: coverage and coherence. The coverage of a story refers to the way it accounts for all the evidence provided. Greater coverage makes the story more acceptable and increases confidence in it. The coherence of a story affects its acceptability and confidence. Coherence consists of three components: consistency, completeness, and plausibility. Consistency means that the story has no internal contradictions. Plausibility means that the story aligns with real or imagined events. Completeness means that the story has all its parts. Together, these components form the coherence of the story.

Furthermore, the story must clearly demonstrate the causal relationships between events. It should show how one event leads to another, establishing a chain of cause and effect that makes the narrative logical and easy to follow.

2.3 Idioms

Pennington and Hastie's model provides a blueprint for how people intuitively construct and understand narratives. These intuitive story structures can be mapped to the formalism of Bayesian networks by using narrative idioms (Vlek et al., 2014). These narrative idioms are a technique that helps modelling by dividing it into smaller, more manageable parts. These idioms represent common patterns of inference used in legal discussions. By doing so, they can illustrate intricate evidence and theories effectively. Additionally, these idioms highlight the importance of relevance and dependence in legal contexts, aligning with the structure of Bayesian networks without focussing on specific numerical probabilities Lagnado et al. (2013).

2.3.1 The scenario idiom

One of the narrative idioms is the scenario idiom. This idiom is designed to model criminal cases by structuring the narrative around a central scenario node. This node connects to all elements of the scenario, ensuring that the evidence supporting any part of the scenario influences the probability of the entire scenario (Vlek et al., 2015). This idiom captures the coherence of a scenario by ensuring that if the scenario node is true, all its connected elements must also be true, effectively modelling the logical flow and dependencies within a scenario. This can be seen in Figure 2.1.

This is formalised by defining a scenario scheme idiom as a Bayesian network fragment. The graph of this fragment consists of a boolean scenario node, denoted as (ScN), which represents the scenario as a whole, and boolean nodes (E), each representing an element of the scenario scheme. The edges of the graph consist of unlabelled connections from (ScN) to each element node (E), shown as double arrows. Furthermore, the graph may contain labelled connections between the element nodes (E_1) and (E_2) with a label (x).

The probability table for each element node (E) is constrained such that for any assignment to the parent nodes within the same scenario, the probability of (E) being true given that (ScN) is true, together with the parent nodes $\mathrm{pa}_S(E_i)$, is equal to 1. Formally, this is expressed as:

$$\Pr(E = T \mid ScN = T, \mathrm{pa}_S(E_i)) = 1 \qquad (2.3)$$

This constraint ensures that if the scenario node is true, each connected element node must also be true with a probability of 1, as illustrated in Figure 2.1. This mechanism maintains the coherence

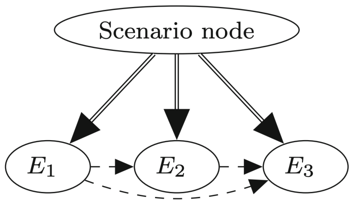

Scenario node	T	F
E_1 =T	1	...
E_1 =F	0	...

Figure 2.1: The scenario idiom, taken from (Vlek et al., 2016), showing how it is connected to the element nodes, and how this influences the CPT.

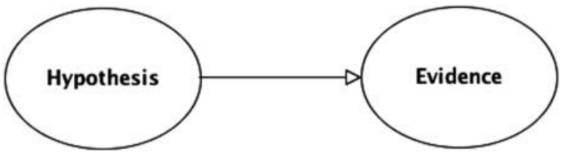

Figure 2.2: The evidential idiom, taken from (Lagnado et al., 2013), showing how evidence is connected to a hypothesis.

and dependency structure among the elements of the scenario. Since the scenario node itself is not directly observed, it effectively manages the logical relationships and dependencies between the elements, ensuring that evidence affecting one part of the scenario influences the entire scenario.

2.3.2 The evidential idiom

In addition to the scenario idiom, the evidential idiom will also be used in this paper. This idiom, introduced by Lagnado et al. (2013), is based on the relationship between evidence and the hypotheses it supports. Each piece of evidence is connected to one or more nodes, as seen in Figure 2.2, which illustrates how the evidence impacts the probability of different hypotheses. This idiom helps to evaluate the strength and relevance of evidence in specific claims.

The relationship between the hypothesis (H) and the evidence (E) in Figure 2.2 indicates that the likelihood of the evidence depends on the hypothesis. Specifically, the probability of observing the evidence given the hypothesis is true, $\Pr(E|H)$, differs from the probability given the hypothesis is false, $\Pr(E|\neg H)$. The strength of this dependency is measured by the likelihood ratio, the quotient of these probabilities:

$$\text{Likelihood Ratio}(LR) = \frac{\Pr(E|H)}{\Pr(E|\neg H)} \tag{2.4}$$

For example, in a fraud case, a financial document showing irregular transactions would be an evidence node connected to the hypothesis node representing the occurrence of fraud. The connection indicates how the presence of the document influences the probability of fraud, allowing the network to update beliefs based on the strength of the evidence.

2.3.3 Connections

Nodes are connected to each other in the network by directed edges. There are two types of connections between the hypothesis nodes, temporal and causal (Vlek et al., 2015).

Temporal connections in Bayesian networks represent the sequence of events over time. They describe the order in which events occur without implying that one event causes the next. In the case of $A \to B \to C$, A happens before B, and B happens before C, but they do not imply that A causes B or B causes C.

Causal connections, on the other hand, indicate a cause-and-effect relationship between events or variables. In a Bayesian network, a node X is said to causally influence node Y if changes in X directly affect the probability of Y. Although this might seem similar to conditional probabilities, it is important to note that causal influence implies a direct effect. In contrast, conditional probabilities refer to the likelihood of one event given another, without necessarily implying direct causation. These connections are important in understanding how different evidence, hypotheses, and scenarios interact and influence each other.

2.4 Legal framework

Passive official corruption involves the acceptance or solicitation of gifts by a public official in circumstances where it is known or should be reasonably suspected that these gifts are intended to influence their official duties Rechtbank Rotterdam (2018b). The concept includes not only direct exchanges of gifts for specific actions, but also situations where gifts are given to foster a relationship that may result in preferential treatment. This legal framework, based on the Dutch book of criminal law, article 262/263, emphasises the importance of officials recognising the implications of accepting gifts and avoiding actions that could be perceived as corrupt, even if there is no explicit exchange of favours.

3 Methods

Building on the model as explained in the previous chapter, this study applies this to a different type of criminal case. The following section outlines the selected case, how the model was created, and how the evidence strength was found.

3.1 The case

The Rotterdam District Court sentenced six people for corruption on February 22, 2018, in the Dotterbloem case (Rechtbank Rotterdam, 2018a). The sentences ranged from community service to 12 months in prison for the main defendant. The investigation, which began in December 2012 following an anonymous tip, focused on corruption involving vehicle procurement for the Dutch police and the Defence Ministry. Eleven people, including six civil servants and five civilians, were accused, but five were acquitted.

The main defendant, a 65-year-old mobility policy officer at the Ministry of Defence, was found guilty of passive bribery (*passieve ambtelijke omkoping*) and violating official secrets (*schending ambtsgeheim*). He leased vehicles under non-market conditions, used fuel cards for personal vehicles, and accepted non-work-related trips from the auto industry. His unique position and access to confidential information were misused for personal gain, leading to a 12-month prison sentence. Other convicted civil servants were deemed naively complicit, failing to recognise the impropriety of accepting personal benefits linked to their official roles.

In order to analyse this extensive and complex case, this paper focusses on one offence by one person. The original text of the case can be found at (Rechtbank Rotterdam, 2018b).

3.2 Creating the network

According to Vlek et al. (2015), the following four steps will be needed to create the network: collect, unfold, merge, and include.

Hugin, a specialised software, was used to construct the Bayesian network. Designed for building and analysing Bayesian networks, it offers an interface for model creation and probabilistic inference. Despite its wide range of functionalities, its features can be complex and involve a learning curve.

3.2.1 Collect

The first step is to collect all relevant scenarios from the case. For this, the positions of both the defendant and the prosecutor are taken. These two scenarios are the two explained in court. Other scenarios could be thought of, but these are not mentioned in the case and thus are not modelled in this network.

The question that needs to be answered in this case is whether the defendant could reasonably have suspected that the claimed actions performed by the company were done to move him to commit fraudulent actions. According to the position of the defendant, this scenario will be:

> Approval by integrity officer AND leased at market price, THUS the defendant did not need to have reasonable suspicion of fraud.

And that of the prosecution will be:

> There was private use AND it was leased for lower than market price, THUS the defendant should have had reasonable suspicion of fraud.

3.2.2 Unfold

The second step is to unfold the scenarios, which means breaking down each scenario into smaller, more detailed sub-scenarios or elements as necessary.

Initially, scenarios are created using the scenario idiom. For each element within these scenarios, a determination is made whether further unfolding is needed by systematically asking specific questions.

Firstly, it is considered whether there is direct evidence connected to the element node. If such evidence exists, no further unfolding is required. If there is relevant evidence for details related to this element or if it is possible to find such evidence, then further unfolding becomes necessary. This is repeated iteratively for each element until no further unfolding is required. This thorough approach ensures that all elements are either adequately detailed and supported by evidence or broken down into smaller elements.

For each element, except for *Reasonable suspicion*, there is evidence to be included. *Reasonable suspicion* is reliant on the previous nodes and cannot be broken down into smaller elements any more.

Two helper nodes are introduced to reduce the complexity of the *Reasonable suspicion* nodes. This is not part of the scenario and thus purposefully not connected to the scenario node. It has been chosen that if any of the previous nodes are true, this node will be true (thus, with an OR construction and not an AND construction).

When constructing logical relationships in scenario modelling, the decision to use AND or OR constructions depends on the underlying probabilities and the nature of the dependencies between elements. Typically, an AND construction is used by default because it requires all conditions to be true for the scenario to hold, providing a stricter and more conservative approach to model dependencies. However, there are specific cases where an OR construction becomes more appropriate, especially in the context of the 'explaining away' effect (Wellman & Henrion, 1993), where the presence of one cause reduces the likelihood of needing other potential causes to explain an outcome.

For example, if both A and B are potential causes of C, confirming A as the cause of C reduces the necessity to attribute C to B. Therefore, to determine the appropriate construction, consider whether all conditions must be met (AND) or if any single condition is sufficient (OR). Employ an OR construction when the presence of one factor lessens the need for others and an AND construction for scenarios requiring strict verification.

3.2.3 Merge

After the unfolding, the two scenarios are merged. A constraint node is connected to each of the scenario nodes in the collection of scenarios. There is a connection from each scenario node to the constraint node. The constraint node has values corresponding to each scenario and one additional value, NA (not applicable), indicating an illegal combination of nodes. The CPT is designed so that, unless exactly one scenario is true, the constraint node will have the value NA. The constraint node is set to ensure that the prior probabilities of the scenario nodes behave as desired, while setting the probability of NA to 0 ensures that multiple scenarios cannot be true simultaneously, nor can it be that none of the scenarios are true (Fenton et al., 2013).

3.2.4 Include evidence

The final step in building the Bayesian network involves integrating the evidence into the network. For each available piece of evidence, a corresponding node is created and connected to the relevant element node it supports within the scenario.

First, it is crucial to identify the relevant evidence. Each piece of evidence from the case must be examined to determine which element or hypothesis it supports or contradicts. This step includes all relevant information in the network, ensuring a thorough analysis.

Once the relevant evidence is identified, the next step is to create the evidence nodes. For each piece of evidence, a separate node is created within the network. Creating distinct nodes for each piece of evidence is essential for later calculating the strength of each piece of evidence. This allows for the quantification of the influence of each piece of evidence on the overall case.

After creating the evidence nodes, the next task is to establish connections. Each evidence node is connected to the element node it directly supports or influences.

In cases where a piece of evidence can be interpreted in multiple ways, additional nodes are created to represent the different interpretations of the evidence. These interpretation nodes are then connected to the original evidence node. Some evidential pieces might be interpreted differently by the prosecution and defence, requiring a separate interpretation node to capture these varying perspectives. This approach ensures that all possible interpretations of the evidence are considered in the network, maintaining the network's completeness.

3.3 Creating CPTs

Each node in the network will have a conditional probability table (CPT). The values related to how likely it is that the element will be true or false are a simplification of the interpretation of probability phrases, as researched in Willems et al. (2020). Table 3.1 shows the numerical mapping used.

State	Value
NA or False	0
Very unlikely	0.1
Unlikely	0.3
Neutral	0.5
Likely	0.7
Very likely	0.9
True	1

Table 3.1: Numerical mapping of probability phrases

To perform the calculations, Bayes' theorem requires the initial prior probability of the hypothesis $\Pr(H)$, regardless of the evidence. Accurately assessing this prior probability is not straightforward. The presumption of innocence, also present in Dutch law, requires that the prior is set in favour of the defendant (Allen et al., 1994). However, it is unclear what the exact value should be. In this research, it will be assumed that the defendant is "very likely" to be innocent.

Additionally, assigning values to certain other CPTs within the network is also challenging. Creating a CPT without prior information from evidence will mean that assumptions have to be made. This means that all event will be considered unlikely if the scenario it is connected to is not true. For example, CPT for *Private use*, will look as follows:

scenario_prosecutor	F	T
private_use F	0.7	0
private_use T	0.3	1

Table 3.2: CPT for *Private use*

Here, according to the theory of the scenario idiom, if the scenario is true, then this node must also be true. In all other cases, it is unlikely that the defendant used the car for private use, without any evidence present.

3.4 Calculating evidence strength

To assess the strength of evidence, a method was used to analyse the likelihood ratio of observation of evidence. The probability of each scenario given the evidence, denoted as $\Pr(s|e)$, was calculated and compared to the prior probability of the scenario, $\Pr(s)$.

The measure of evidential strength used is the ratio $\Pr(s|e)/\Pr(s)$. A ratio greater than 1 indicates supporting evidence for the scenario, while a ratio less than 1 indicates attacking evidence. This method quantifies the strength of evidence for or against each scenario (Fenton et al., 2014).

For each piece of evidence, this ratio was calculated to determine its impact on the probability of the corresponding scenario. By systematically applying this measure, it is possible to identify and report the strength of supporting or attacking evidence for each scenario.

4 Results

The resulting Bayesian network, constructed according to the method explained above, consists of multiple nodes representing various pieces of evidence, elemental nodes which follow the structure of the story, helper nodes to manage complexity, and scenario nodes. The final network structure is shown in Figure 4.1 and explained in Table 4.1.

4.1 The network

To construct the network, the process as described in the Methods section was followed: collect, unfold, merge, and include evidence.

4.1.1 Collect

The elements of the network were derived from the positions of both the prosecutor and the defendant. The prosecutor's scenario suggests that the defendant should have had reasonable suspicion of fraudulent activities due to the private use of vehicles and leasing them at a lower than market price. On the other hand, the defendant's scenario argues that the lease agreements were approved by an integrity officer and leased at market price, denying any need for reasonable suspicion of fraud. Furthermore, all relevant pieces of evidence

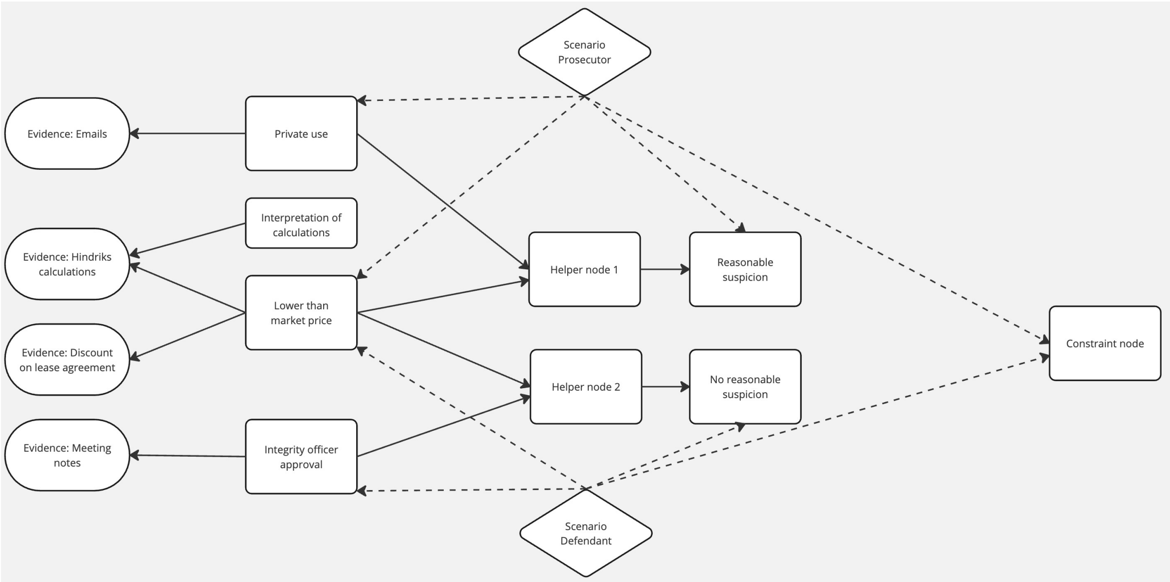

Figure 4.1: The Bayesian network illustrating the chosen segment of the case, showing the scenarios for both the prosecutor and the defendant, together with the presented evidence.

Type	Name	Meaning	Scenario	Evidence
Scenario node	Scenario prosecutor	The prosecutor's scenario	Prosecutor	None
Scenario node	Scenario defendant	The defendant's scenario	Defendant	None
Story element	Private use	If the car was used for private purposes	Prosecutor	Emails
Story element	Lower than market price	If the lease price was below market value	Prosecutor Defendant	Hindriks calculations and Discount on lease agreement
Story element	Reasonable suspision	If there was reasonable suspicion of fraudulent activities.	Prosecutor Defendant	None
Story element	Integrity officer approval	If the lease agreement was approved by the integrity officer	Defendant	Meeting notes
Helper node	Helper node 1 and 2	Used to manage the complexity of relationships between other nodes	None	None
Constraint	Constraint node	Ensures that only one scenariocan be true at a time	Prosecutor Defendant	None

Table 4.1: Detailed description of the nodes in the Bayesian network

were collected from the court documents. This ensured inclusion of all relevant perspectives and evidence from the Dotterbloem case for constructing the Bayesian network.

4.1.2 Unfold

In this step, each scenario was broken down into more detailed sub-scenarios or elements. This process involved deconstructing the broader narratives into smaller and manageable parts.

Initially, the scenario idiom was applied to create high-level scenarios for both the prosecutor and the defendant. Then each element of these scenarios was analysed to determine if further subdivision was necessary. To manage complexity, two helper nodes were introduced for the element *Reasonable suspicion.* These nodes reduced the complexity by combining influences of connected sub-elements. If any contributing factor was true, the relevant helper node would also be true, thus simplifying the network structure.

This resulted in two scenarios: one from the prosecutor's perspective, highlighting the elements of private use, below market lease price, and reasonable suspicion of fraud and another from the defendant's perspective, highlighting the integrity officer's approval, market-rate leasing terms, and the lack of reasonable suspicion.

4.1.3 Merge

In this step, the two detailed scenarios of the prosecutor and the defendant were combined into a single Bayesian network. A constraint node was introduced to ensure that only one scenario could be true at a time, maintaining the logical consistency of the network. This node was connected to each scenario node, enforcing a mutual exclusivity constraint. Each scenario node was then connected to its relevant elements.

The edges in the network represent the conditional dependencies between the nodes, which can be temporal or causal. In this research, edges are not specifically labelled, as the focus is on the overall structure and functionality of the network rather than the precise categorisation of each edge.

4.1.4 Include evidence

Lastly, the following pieces of evidence were included into the network:

- *Emails*: There was an email correspondence stating that employees of the car company involved knew that the defendant wanted to lease the car for private use.
- *Hindriks calculations*: This node represents the interpretation of financial calculations on the lease price of the vehicle. The defence pointed to the calculations by the expert Hindriks, which was based on the investment value provided by the owner/director. Hindriks' findings suggest that the lease price was calculated fairly and aligned with market rates, supporting the defence's claim that the leasing terms were legitimate and did not indicate fraud. The prosecutor has interpreted these calculations differently, which is beneficial for their narrative. To model this, an extra node is created, called *Interpretation of calculations.*
- *Discount on lease agreement*: This evidence piece demonstrates that the company provided significant discounts on the delivery of the vehicle, thus contributing to the *Lower than market price* node
- *Meeting notes*: These meeting notes were used to validate the approval of the integrity officer.

4.2 Creating CPTs

Each node is associated with a CPT, which shows how the node relies on its parent nodes. All CPTs can be found in Appendix A. Particularly interesting nodes are the *Lower than market price* node and the *Reasonable suspicion* nodes, which presents a challenge due to the conflicting scenarios proposed by the prosecution and the defence. The complexities and resolutions of this node are discussed in detail in the Discussion section.

4.3 Calculating evidence strength

The strengths of individual pieces of evidence were calculated based on $\Pr(s|e)/\Pr(s)$, where $\Pr(s)$ is the prior probability of the scenario and $\Pr(s|e)$ is the conditional probability given the evidence. A strength greater than one indicates that it is supporting the scenario. A strength between one and zero indicates that the evidence is attacking the scenario.

- *Emails:* The conditional probability of the prosecutor's scenario given the emails is $\Pr(s|emails) = 0.2273$. This results in an evidence strength of 2.273. The probability of the defendant's scenario given the emails is $\Pr(\neg s|emails) = 0.7727$. The strength of this evidence for the defendant's scenario is 0.859. This makes this piece of evidence supporting for the prosecutor's scenario.

- *Discount on lease agreement:* The evidence from the discount on the lease agreement significantly supports the prosecutor's scenario with $\Pr(s|discount) = 0.5$. The strength of this evidence is 5.0, suggesting a five-fold increase in the likelihood of the scenario, making it a very influential piece of evidence. The probability of the defendant's scenario given the discount is $\Pr(\neg s|discount) = 0.5$. The strength of this evidence for the defendant's scenario is 0.556. This makes this piece of evidence supporting for the prosecutor's scenario.

- *Hindriks calculations (interpretation of the defence):* The conditional probability of the prosecutor's scenario given the Hindriks calculations, given it is interpreted as the defence proposes, is $\Pr(s|Hindriks\ calculations\ (defence)) = 0.0122$. This results in an evidence strength of 0.122. The probability of the defendant's scenario given this interpretation of the Hindriks calculations is $\Pr(\neg s|Hindriks\ calculations\ (defence)) = 0.9878$. The strength of this evidence for the defendant's scenario is 1.098. This makes this piece of evidence supporting for the defendant's scenario.

- *Hindriks calculations (interpretation of the prosecutor):* The conditional probability of the prosecutor's scenario given the Hindriks calculations, given it is interpreted as the prosecutor proposes, is $\Pr(s|Hindriks\ calculations\ (prosecutor)) = 0.5$. The probability of the defendant's scenario given this interpretation of the Hindriks calculations is $\Pr(\neg s|Hindriks\ calculations\ (prosecutor)) = 0.5$. The strength of this evidence for the defendant's scenario is 0.556.

- *Documentation:* The documentation's conditional probability is $\Pr(s|documentation) = 0.0753$. This yields a strength of 0.753, indicating that this piece of evidence slightly weakens the prosecutor's scenario when considered in isolation. The probability of the defendant's scenario given the documentation is $\Pr(\neg s|documentation) = 0.9247$. The strength of this evidence for the defendant's scenario is 1.027.

The piece of evidence with the most strength for the prosecutor's scenario is *Discount on the lease agreement*. This evidence significantly supports the prosecutor's scenario with a strength of 5.0, indicating a five-fold increase in the likelihood of the scenario compared to the prior probability. The addition of *Hindriks calculations* also provides a significant strength of 5.0, indicating that both pieces of evidence strongly support the scenario of the prosecutor. The piece of evidence with the most strength for the defendant's scenario is *Hindriks calculations (interpretation of the defence)*. This evidence slightly supports the defendant's scenario with a strength of 1.098, indicating that it increases the likelihood of the defendant's scenario slightly, compared to the prior probability.

5 Discussion

In this research, a Bayesian network has been applied to model criminal fraud cases. This section will focus on the benefits and limitations of integrating the scenario idiom in this context. Combining Bayesian networks with scenario schemes offers a structured approach to understanding complex legal cases. Scenario idioms help in finding out if all elements of the scenario are consistent and thoroughly documented, helping in transforming the network structure into a clear narrative.

5.1 Benefits

A major benefit of this approach is that it enables judges and juries to make decisions through a scenario-based method while including probabilistic data. This framework improves comprehension by helping decision makers understand the content and implications of various pieces of evidence. In this case, for example, this approach facilitated the evaluation of complex financial documents, email correspondences, and even multiple interpretations of one piece of evidence.

One specific benefit observed was the ability to quantify the strength of evidence. For example, the analysis of emails indicating private vehicle use precisely quantified how this evidence supported the prosecution's scenario. These numerical values provided a clear measure of how strongly the evidence supported the scenario.

Lastly, the use of scenario schemes provided a structured way to assess the coherence of the narratives presented by both the prosecution and the defence. This method ensured that all elements of the scenarios were accurately and thoroughly represented, allowing for a more transparent evaluation of the case.

5.2 The scenario node

However, there are also notable limitations. Although scenarios provide valuable context, it is not always necessary to attach a scenario to every node. Some nodes represent elements that do not require additional narrative context. In this research, this was shown by the helper nodes. As these were not part of the narrative, they do not need the scenario node as their parent node. This could lead to unnecessary complexity. Moreover, merging two distinct storylines into one Bayesian network adds considerable complexity, especially when building CPTs. The CPTs that result can be challenging to handle, requiring intricate construction and validation. Additionally, most cases, either fraud or not, involve more than two scenarios. The Dotterbloem case in the study was simplified, but more intricate cases might encounter numerous complexity issues.

5.3 Integration of scenarios

The modelling of certain elements presented significant semantic challenges, especially when integrating the two scenarios within the Bayesian network. The *Lower than market price* node and the nodes that represent reasonable suspicion were particularly problematic.

One particularly interesting node in the Bayesian network is the *Lower than market price* node. This node presents a challenge due to the conflicting scenarios proposed by the prosecution and the defense. According to the CPT for this node, if the prosecution's scenario is true while the defense's scenario is false, then the *Lower than market price* node must indeed be true. This aligns with the prosecution's argument that the car was rented for less than the market value.

However, the situation becomes more complex when the defence scenario is true and the prosecution scenario is false. Semantically, to support the defence's claim that the car was not rented below the market price, this node should be false. This creates a contradiction with the theory of the scenario idiom, which says that if a scenario is true, all its associated nodes should also be true. According to this theory, for the defence's scenario to be true, the *Lower than market price* node would incorrectly have to be true, implying that the car was rented below market price.

To resolve this conflict and accurately model the network, the semantic meaning of the element was prioritised over the rigid application of the scenario idiom theory. This adjustment ensures that the network faithfully represents the defence scenario, in which the car was not rented for less than the market price, maintaining logical consistency.

Since this node is connected to both scenario nodes, it functions as a constraint node. In a constraint node, this situation is represented as "NA". However, in this particular case, applying such a constraint is not applicable to a normal element node. The constraint node in the network will ensure that these two combinations (both scenarios being true and both scenarios being false) will not occur.

The concept of reasonable suspicion was a similar challenge, but with a different solution. Originally represented by a single node, it was split into *Reasonable suspicion* and *No reasonable suspicion* nodes to avoid conflicts like with the *Lower than market price* node, and to represent the defence and prosecution scenarios clearly. This solves the issues that showed for the *Lower than market price* node, but the nodes do lose important semantic information. This is because they do not have both scenarios as their parent nodes, but one of the two.

These adjustments revealed a limitation in the network's ability to integrate elements from both scenarios without contradictions. This limitation shows the complexity of modelling legal cases with Bayesian networks, where shared elements must accurately reflect their roles in both narratives.

5.4 Responsibilities

Addressing the inclusion of story elements, it was noted that the final decision on the relevance of the nodes usually lies with the decision maker, not the person constructing the network. This distinction is crucial, as it highlights the role of the legal decision maker in interpreting the network's output and ensuring that only pertinent information influences the judgment. The network must be designed to accurately accommodate issues such as those shown in this research. These issues are very important in the construction of the network and addressing them compels the network builder to make decisions that ideally should not be their responsibility.

5.5 The problem of priors

The problem of establishing priors in Bayesian networks remains a significant issue. The presumption of innocence dictates that the prior probability should favour the defendant. However, determining the exact value of this prior is not straightforward and can significantly influence the outcomes of the network.

5.6 The variation idiom

Lastly, the potential application of the variation idiom, which was not used in this research, suggests an area for future exploration. The variation idiom could offer a structured way to handle different interpretations of the same evidence, thus enriching the network's ability to model complex legal scenarios more accurately.

6 Conclusion

This thesis investigated the application of Bayesian networks, incorporating scenario idioms, to model criminal fraud cases. The study focused on the Dotterbloem case to examine how well these methods, previously used in murder cases, translate to the domain of fraud. This analysis revealed several benefits and limitations emerged.

One of the significant benefits observed was the ability of the method to accommodate various forms of evidence. This capability is particularly important in fraud cases, which often involve interpreting financial documents, emails, and other circumstantial evidence. The scenario idioms facilitated the construction of a narrative that linked these various pieces of evidence into a coherent whole, assisting legal decision makers in their assessments.

However, the study also identified several limitations. Fraud cases often require more complex Bayesian networks, which can be challenging to construct and validate. The merging of different stories into a single network introduced significant semantic challenges, such as nodes that must be false according to one scenario but true according to another, and the appearance where nodes function as constraint nodes when they should not.

In addition, the challenge of establishing accurate prior probabilities within the network remains an issue. The presumption of innocence requires setting priors that favour the defendant, but determining the exact value of these priors is not straightforward and can significantly impact the outcomes of the network.

Future research should continue to improve these methods, address their limitations, and explore their application to other types of legal cases.

In conclusion, the adaptation of Bayesian networks with scenario idioms for fraud cases might offer an approach to legal evidence analysis. Although scenario idioms offer structure, they require careful design to ensure semantic and logical consistency, crucial for reliable and transparent evidence interpretation.

References

Allen, R. J., Balding, D. J., Donnelly, P., & Friedman, R. (1994). Probability and proof in State v. Skipper: an internet exchange. *Jurimetrics J.*, *35*, 277.

Dawid, A. P. (2010). Beware of the DAG! In *Causality: objectives and assessment* (pp. 59–86).

Desai, S. C., Reimers, S., & Lagnado, D. A. (2016). Consistency and credibility in legal reasoning: A Bayesian network approach. In *Cogsci.*

Fenton, N., Berger, D., Lagnado, D., Neil, M., & Hsu, A. (2014). When 'neutral' evidence still has probative value (with implications from the Barry George Case). *Science & Justice*, *54*(4), 274–287.

Fenton, N., Neil, M., & Lagnado, D. A. (2013). A general structure for legal arguments about evidence using Bayesian networks. *Cognitive science*, *37*(1), 61–102.

Hans, V. P., & Saks, M. J. (2018). Improving judge & jury evaluation of scientific evidence. *Daedalus*, *147*(4), 164–180.

Jensen, F. V., & Nielsen, T. D. (2007). *Bayesian networks and decision graphs* (Vol. 2). Springer.

Kammen, & Moudy. (2023, Nov). *How does someone prove intent in a fraud case?* Retrieved from `www.kammenlaw.com`

Lagnado, D. A., Fenton, N., & Neil, M. (2013). Legal idioms: a framework for evidential reasoning. *Argument & Computation*, *4*(1), 46–63.

Pearl, J. (1988). Embracing causality in default reasoning. *Artificial Intelligence*, *35*(2), 259–271.

Pennington, N., & Hastie, R. (1991). A cognitive theory of juror decision making: The story model. *Cardozo L. Rev.*, *13*, 519.

Rechtbank Rotterdam. (2018a, 2). Celstraf en taakstraffen voor corruptie bij levering voertuigen politie en Defensie. Retrieved from `https://shorturl.at/5cNrz`

Rechtbank Rotterdam. (2018b, 2). *Uitspraak.* Retrieved from `https://shorturl.at/iYCSL`

Vlek, Prakken, H., Renooij, S., & Verheij, B. (2014). Building Bayesian networks for legal evidence with narratives: a case study evaluation. *Artificial intelligence and law*, *22*, 375–421.

Vlek, Prakken, H., Renooij, S., & Verheij, B. (2015). Constructing and understanding Bayesian networks for legal evidence with scenario schemes. , 128–137.

Vlek, Prakken, H., Renooij, S., & Verheij, B. (2016). A method for explaining Bayesian networks for legal evidence with scenarios. *Artificial Intelligence and Law*, *24*, 285–324.

Wellman, M. P., & Henrion, M. (1993). Explaining 'explaining away'. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *15*(3), 287–292.

Willems, S., Albers, C., & Smeets, I. (2020). Variability in the interpretation of probability phrases used in Dutch news articles — a risk for miscommunication. *Journal of Science Communication*, *19*(2), A03.

A Conditional probability tables

A.1 Scenarios

The truth tables of the scenario nodes (A.1 and A.2) represent the priors of the networks. The main issue of this representation is that there is no prior available. As explained in Chapter 3, a presumption of innocence resulted in the following CPTs:

state	**value**
false	0.9
true	0.1

Table A.1: CPT for *Scenario prosecutor*

state	**value**
false	0.1
true	0.9

Table A.2: CPT for *Scenario defendant*

When the two scenarios are merged, a constraint node is connected to each of the scenario nodes. There is a connection from each scenario node to the constraint node. The constraint node has values corresponding to each scenario and one additional value, NA (not applicable), indicating an illegal combination of nodes. The CPT is designed so that, unless exactly one scenario is true, the constraint node will have the value NA. The constraint node is set to ensure that the prior probabilities of the scenario nodes behave as desired, while setting the probability of NA to 0 ensures that multiple scenarios cannot be true simultaneously, nor can it be that none of the scenarios are true Fenton et al. (2013).

scenario_defendant	F		T	
scenario_prosecutor	F	T	F	T
scenario_prosecutor T	0	1	0	0
scenario_defendant T	0	0	1	0
NA	1	0	0	1

Table A.3: CPT for *Constraint node*

A.2 Elemental nodes

For the *Private use* node (A.4), if the scenario node is true, this node must also be true. On the other hand, if the scenario node is false, it remains possible that the car was used for private purposes, but this is considered unlikely.

scenario_prosecutor	F	T
private_use F	0.7	0
private_use T	0.3	1

Table A.4: CPT for *Private use*

Similarly, for the *Integrity officer approval* node (A.5), if the scenario node is true, then the *Integrity officer approval* node must also be true. If the scenario node is false, it is still possible that the lease agreement was approved by an integrity officer. This is considered to be likely.

scenario_defendant	F	T
integrity_officer_approval F	0.3	0
integrity_officer_approval T	0.7	1

Table A.5: CPT for *Integrity officer approval*

The *Lower than market price* node presents a challenge due to the conflicting scenarios proposed by the prosecution and the defence. According to the CPT for this node (A.6), if the prosecution's scenario is true while the defence's scenario is false, then the *Lower than market price* node must indeed be true. This aligns with the prosecution's argument that the car was rented for less than the market value.

However, the situation becomes more complex when the defence scenario is true and the prosecution scenario is false. Semantically, to support the defence's claim that the car was not rented below the market price, this node should be false. This creates a contradiction with the theory of the scenario idiom, which says that if a scenario is true, all its associated nodes should also be true. According to this theory, for the defence's scenario to be true, the *Lower than market price* node would incorrectly have to be true, implying the car was rented below market price.

To resolve this conflict and accurately model the network, we prioritise the semantic meaning of the element over the rigid application of the scenario idiom theory. This adjustment ensures that the network faithfully represents the defence scenario, in which the car was not rented for less than the market price, maintaining logical consistency.

Since this node is connected to both scenario nodes, it functions as a constraint node. This means that, in theory, it enforces a rule where multiple scenarios cannot be true simultaneously, nor can it allow a situation where none of the scenarios are true. Typically, in a constraint node, this situation is represented as "NA". However, in this particular case, applying such a constraint is not feasible. The constraint node in the network will ensure that these two combinations (both scenarios being true and both scenarios being false) will not occur.

scenario_defendant	F		T	
scenario_prosecutor	F	T	F	T
lower_than_market_price F	0.7	0	1	0
lower_than_market_price T	0.3	1	0	1

Table A.6: CPT for *Lower than market price*

The CPTs for the two helper nodes use an OR structure, which simplifies the complexity of certain nodes by ensuring that if any contributing factor is true, the helper node itself will be true. This structure is based on the principle expressed in Equation A.1

$$\frac{Pr(C=T|A=T,B=T)}{Pr(C=T|A=T,B=F)} \leq \frac{Pr(C=T|A=F,B=T)}{Pr(C=T|A=F,B=F)} \tag{A.1}$$

For the first helper node A.7, if either the *Private use* node or the *Lower than market price* node is true, the helper node will also be true. This is represented in the CPT as follows: If both contributing nodes are false, the helper node is false; if either one or both of the contributing nodes are true, the helper node is true. This OR structure is useful in scenarios where the truth of any single element justifies the truth of the overall helper node.

private_use	F		T	
market_price	F	T	F	T
helper_node1 F	1	0	0	0
helper_node1 T	0	1	1	1

Table A.7: CPT for *Helper node 1*

Similarly, the second helper node aggregates the influences of multiple nodes. For example, if the *Integrity officer approval* node is true, the second helper node will be true. The CPT for this node also follows the OR structure: It is false only if all contributing nodes are false; it is true if any one of the contributing nodes is true. This approach ensures that the presence of any significant evidence supports the helper node, simplifying the complexity of the network while maintaining logical consistency.

As explained in the method section, the node representing whether there should have been reasonable suspicion of fraud is split into two nodes.

The node representing reasonable suspicion is structured as follows:

scenario_prosecutor	F		T	
helper_node1	F	T	F	T
reasonable_suspicion F	0.9	0.1	0	0
reasonable_suspicion T	0.1	0.9	1	1

Table A.8: CPT for *Reasonable suspicion*

If the prosecution's scenario is false and the helper node is also false, it is very likely that there is no reasonable suspicion.

Similarly, the node representing no reasonable suspicion follows the same structure, but the interpretation is inverted. If the defendant's scenario is false and the helper node is also false, it is very likely that the defendant was not obligatory to have reasonable suspicion at the situation.

scenario_defendant	F		T	
helper_node2	F	T	F	T
no_reasonable_suspicion F	0.9	0.1	0	0
no_reasonable_suspicion T	0.1	0.9	1	1

Table A.9: CPT for *No reasonable suspicion*

An additional node in the network is the *Interpretation of calculations* node (A.10). This node captures the ambiguity in how financial calculations are interpreted within the context of the case. The node influences the evidence node and is defined with equal probabilities for being true or false:

state	**value**
defence	1
prosecutor	1

Table A.10: CPT for *Interpretation of calculations*

A.3 Evidence nodes

The CPT for the evidence node *Documentation* represents the relationship between the presence of documentation and the approval by an integrity officer. If the integrity officer approval is false, the probability of not having documentation is very likely (0.9), while the probability of having documentation is very unlikely (0.1). This aligns with the expectation that in the absence of approval, supporting documentation is likely to be absent. If the integrity officer approval is true, the probability of having documentation is very likely (0.9), and the probability of not having documentation is very unlikely (0.1). This reflects the expectation that when approval is granted, it is usually documented.

The CPT for *Discount on lease agreement* represents the relationship between the presence of a discount on the lease agreement and whether the lease price is lower than the market price. If the lease price is not lower than the market price (false), the probability of not having a discount on the lease agreement is very likely (0.9), while the probability of having a discount is very unlikely (0.1). This aligns with the expectation that without a discount, the lease price is likely to be at or above the market rate. If the lease price is lower than the market price (true), the probability of having a discount on the lease agreement is very likely (0.9), and the probability of not having a discount is very unlikely (0.1). This reflects the expectation that a lease price below the market rate is usually associated with some form of discount.

integrity officer approval	F	T
documentation F	0.9	0.1
documentation T	0.1	0.9

Table A.11: CPT for *Documentation*

lower than market price	F	T
discount_on_lease_agreement F	0.9	0.1
discount_on_lease_agreement T	0.1	0.9

Table A.12: CPT for *Discount on lease agreement*

The CPT for *Emails* represents the relationship between the presence of emails indicating private use of the vehicle and whether the vehicle was indeed used for private purposes. If the vehicle was not used for private purposes (false), the probability of not having emails indicating private use is very likely (0.9), while the probability of having such emails is very unlikely (0.1). This aligns with the expectation that in the absence of private use, there would likely be no emails suggesting otherwise. If the vehicle was used for private purposes (true), the probability of having emails indicating private use is very likely (0.9), and the probability of not having such emails is very unlikely (0.1). This reflects the expectation that private use would be documented or discussed in emails.

The evidence table for *Interpretation of calculations* is structured to represent the relationship between the interpretation of financial calculations and whether the lease price is considered lower than the market price, as viewed by both the defence and the prosecution. This CPT indicates that the interpretation of the financial calculations is heavily influenced by the perspectives of both the defence and the prosecution. The interpretation of each support their respective arguments whether the lease price is lower than the market price.

From the defence perspective, if the lease price is not lower than the market price (false), it is very likely (0.9) that the interpretation calculations will align with this view and it is very unlikely (0.1) that they will indicate otherwise. If the lease price is lower than the market price (true), the probability that interpretation calculations align with this view is very unlikely (0.1), and the probability that they indicate that the lease price is lower than the market price is very likely (0.9).

For the prosecution's perspective, if the lease price is not lower than the market price (false), it is very unlikely (0.1) that the interpretation calculations will align with this view and it is very likely (0.9) that they will indicate otherwise. If the lease price is lower than the market price (true), the probability that interpretation calculations align with this view is very likely (0.9), and the probability that they indicate that the lease price is not lower than the market price is very unlikely (0.1).

private use	F	T
emails F	0.9	0.1
emails T	0.1	0.9

Table A.13: CPT for *Emails*

interpretation calculations	defense		prosecutor	
lower than market price	F	T	F	T
hindriks_calculations F	0.1	0.9	0.9	0.1
hindriks_calculations T	0.9	0.1	0.1	0.9

Table A.14: CPT for *Hindriks calculations*