

COMBINING UNCERTAINTIES IN A COURT OF LAW USING BAYESIAN NETWORKS

MA Muller
MSc PhD
Department of Mathematics
University of Stellenbosch

SUMMARY

People generally have difficulty dealing with the counter-intuitive notion of probability, and therefore they often misunderstand aspects of uncertainty. This is particularly significant in a court of law when for example an estimate of the probability of the evidence gets confused with an estimate of the probability of guilt. The circumstantial evidence is especially prone to being handled incorrectly. Professor Fenton at the Queen Mary University of London said, "You could argue that virtually every case with circumstantial evidence is ripe for being improved by Bayesian arguments".¹ In this paper, the evidence in a famous court case is revisited in the context of Bayesian networks.

1 INTRODUCTION

This paper is a sequel of two previous articles² in which comparatively simple cases were considered where Bayesian networks could have been used to obtain an estimate of the (say) probability³ of guilt in a court of law.

For example, at a major athletics championship, a randomly selected athlete is tested for the use of banned substances. The outcome of the test is positive. It is known that the test gives a true reflection of the actual state of affairs with probability 95%.⁴ Therefore, the test is quite accurate, but not

¹ Saini "A Formula for Justice" 2 October 2011 *The Guardian* www.theguardian.com/law/2011/oct/02/formula-justice-bayes-theorem-miscarriage (accessed 2016-07-28).

² Muller "Handling Uncertainty in a Court of Law" 2012 23(3) *Stell LR* 599 600–601 <http://scholar.sun.ac.za/handle/10019.1/98245> (accessed 2016-07-28); Muller "Underestimating the Probability of Coincidence" 2014 35(2) *Obiter* 173 173–187 <http://scholar.sun.ac.za/handle/10019.1/98247> (accessed 2016-07-28).

³ The probability $P(A)$ of an event A is a measure of the strength of one's conviction of the truth that the event A occurs. It is always a number between 0 and 1, or equivalently, a percentage between 0% and 100%; Ross *A First Course in Probability* (2010) 22–57.

⁴ This means that if the athlete indeed took a banned substance, the test would indicate such with probability 95%, and the test would indicate that the athlete did not take a banned substance with probability 5%. Similarly, if the athlete indeed did not take a banned substance, the test would indicate such with probability 95%, and the test would indicate that the athlete did take a banned substance with probability 5%.

perfect. It is also known that 1% of the athletes take banned substances. Surprisingly, it turns out that the probability that this athlete indeed took a banned substance is in fact 16.1%. Clearly, the numerical value of the probability of guilt is widely different from the numerical value of the probability of the evidence. This result was obtained by simple manual calculation from first principles (that is, with “pencil and paper”).⁵ The same result could also be obtained by using a Bayesian network.

Let E and G be the following events:

E : The evidence is presented that the athlete’s test for banned substances is positive.

G : The athlete is guilty.

The following diagram describes the type of Bayesian network we would consider:



This model describes the relationship between events E and G . However, since we already obtained a method for solving the problem from first principles when a single piece of evidence is presented, we prefer not to use Bayesian networks in such cases. However, court cases tend not to be all that simple. The real question is how multiple pieces of evidence of different degrees of uncertainty should be handled. The first step is finding a way of handling *two* contributing pieces of evidence, instead of *one* piece of evidence E as in the above example.

Aspects of the following example may turn up in a court of law if a client who suffered heart failure claims compensation from manufacturers of tobacco products or certain types of fast food. It illustrates how in general the combining of probabilities should be handled in situations that are more complex. We use methods demonstrated in this section in the latter part of the paper.

Example 1

Given a large group of people. We are given the fact that 20% of this population follow unhealthy diets and 10% are smokers. Take any randomly selected person from this group. Let D , S and H be the following events:

D : This person follows an unhealthy diet.

S : This person smokes.

⁵ Muller 2012 23(3) *Stell LR* 599 Examples 5 and 7.


H: This person suffers heart failure.

We have $P(D) = 0.2$ and $P(S) = 0.1$ and we are also given the following table of conditional probabilities:⁶

	Healthy diet ($\sim D$)		Unhealthy diet (D)	
	Does not smoke ($\sim S$)	Smokes (S)	Does not smoke ($\sim S$)	Smokes (S)
No heart failure ($\sim H$)	0.7 (= $P(\sim H (\sim D) \cap (\sim S))$)	0.45 (= $P(\sim H (\sim D) \cap S)$)	0.4 (= $P(\sim H D \cap (\sim S))$)	0.25 (= $P(\sim H D \cap S)$)
Heart failure (H)	0.3 (= $P(H (\sim D) \cap (\sim S))$)	0.55 (= $P(H (\sim D) \cap S)$)	0.6 (= $P(H D \cap (\sim S))$)	0.75 (= $P(H D \cap S)$)

Figure 1

We consider the above by means of a so-called Bayesian network.⁷ In this case, the Bayesian network is described by the following diagram:



The model represents the relationships between events. A line between two nodes suggests that one node has a causal influence on the other. Moreover, it represents these relationships mathematically since each node has associated with it a node probability table.⁸ Figure 1 is the node probability table of the event *H*.

Since the events *D* and *S* are mutually independent events, there is no direct connection between these nodes in the diagram. The probabilities $P(D) = 0.2$ and $P(S) = 0.1$ form our initial understanding of the state of affairs before anything else relating to *H* was known. The node probability table of *H* contains conditional probabilities of the type where the probability of the event *H* appears under the assumption that information about the parent events *D* and *S* are given. These probabilities could, for instance represent known mathematical facts, the results of research, the opinion of an expert, or the opinion of a judge in a court of law.⁹ We may associate this

⁶ If *A* and *B* are any events then $P(B|A)$ is called the conditional probability of *B*, given *A* (that is, the probability of the event *B*, given that event *A* occurs.) The second entry in the bracket is the condition under which the probability of the first entry is given.

$\sim A$ denotes the event when event *A* does *not* occur. For any event *A* we always have $P(A) + P(\sim A) = 1$ (also written $P(A) + P(\sim A) = 100\%$).

$A \cap B$ denotes the event when both events *A* and *B* occur. Clearly $A \cap B = B \cap A$. Sometimes $A \cap B$ is written “*A*&*B*” or “*A* and *B*”.

⁷ Fenton and Neil *Risk Assessment and Decision Analysis with Bayesian Networks* (2013) 69–266; Korb and Nicholson *Bayesian Artificial Intelligence* (2004) 3–146.

⁸ Agena Ltd *Getting Started with AgenaRisk* (2013) 5.

⁹ 12–13.

information with the action of moving down the diagram from D or S to H . The numerical values of similar probabilities associated with moving upwards from H to D or S are as yet unknown. For example, in the appendix the numerical values of the following probabilities are calculated manually from first principles:

- (a) $P(H)$, the probability that a randomly selected person from this group suffers heart failure.
- (b) $P(S|H)$, the conditional probability that a randomly selected person from this group who suffered heart failure is a smoker.
- (c) $P(D \cup S)$,¹⁰ the probability that a randomly selected person from this group indulges in at least one of the following: an unhealthy diet (D), or smoking (S).

If more than two causes of heart failure were given, the calculation would be approached in a similar way as in Example 1. However, the manual approach soon becomes very cumbersome. We prefer an easier way that avoids the manual manipulation. Fortunately, computer software is available that enables us to deal easily with several contributing events. The programmes *AgenaRisk* and *Hugin* are generally aimed at the needs of large companies and corporations, but the user-friendly “lite” versions are adequate for our purposes and could be downloaded from the Internet.¹¹

The following actions are typically undertaken when using such software. First, the given information (that is, the fact that $P(D) = 0.2$, $P(S) = 0.1$ and the node probability table in figure 1) are loaded into the programme. A user manual shows where and how it is done.¹² Then the nodes are linked up with appropriate arrows as in figure 2. Lastly, the programme is activated and the result appears almost immediately. We see below that the computer’s results are the same as those obtained in the appendix.

We draw attention to the fact that the manual calculations in the appendix are merely listed to demonstrate that the same results could just as well have been arrived at by using a Bayesian network on a computer. Tedious Bayesian arithmetic can be almost completely automated.¹³ In practical terms, the approach in the appendix is rarely used. Fenton and Neil conclude a clearly formulated argument about this by stating

“there should be no more need to explain the Bayesian calculations in a complex argument than there should be any need to explain the thousands of circuit level calculations used by a calculator to compute a long division”.¹⁴

¹⁰ If A and B are any events then $A \cup B$ denotes the event when at least one of the events A or B occur. Clearly $A \cup B = B \cup A$. Sometimes $A \cup B$ is written “ A or B ”.

¹¹ www.agenarisk.com; www.hugin.com (accessed 2016-07-28).

¹² Agena Ltd *Getting Started with AgenaRisk* 5.

¹³ Edwards “Influence Diagrams, Bayesian Imperialism, and the Collins Case: An appeal to Reason” 1991 13 *Cardozo LR* 1025 1033.

¹⁴ Fenton and Neil “Avoiding Legal Fallacies in Practice Using Bayesian Networks” 2011 36 *Australian Journal of Legal Philosophy* 114 149.

(a) We determine $P(H)$, the probability that a randomly selected person from this group suffers heart failure.

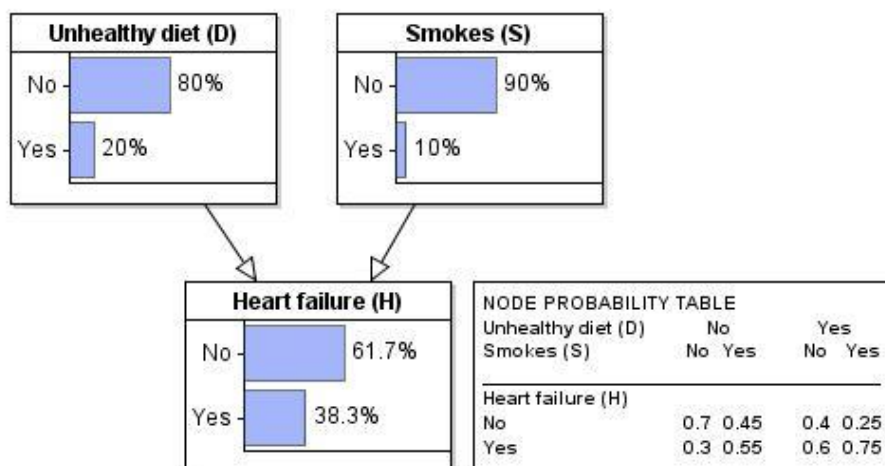


Figure 2

So by the computer-aided method, $P(H) = 38.3\%$.

(b) We determine $P(S|H)$.

AgenaRisk has a facility that calculates the conditions under which a particular observation would follow. Given that a randomly selected person from this group indeed had heart failure (that is, given $P(H) = 100\%$) and given the information in figure 2, then figure 3 gives the probability that he smokes.

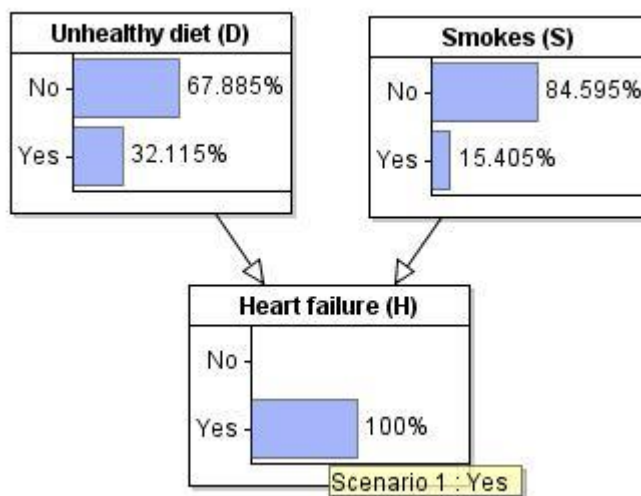


Figure 3

By the computer-aided method, $P(S|H) = 15.4\%$.

- (c) We determine $P(DUS)$, the probability that a randomly selected person from this group indulges in at least one of the following: an unhealthy diet (D), or smoking (S).

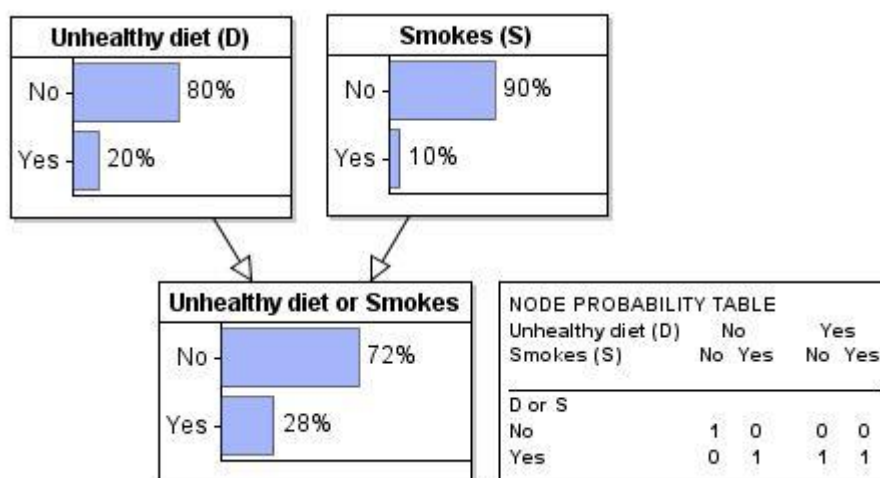


Figure 4

By the computer-aided method, $P(DUS) = 28\%$.

Since the probabilities in figure 1 could be based on people's observations or opinions, we may wish to see what the effect of different sets of input values (other than the numerical values in figure 1) would be. The computer-aided method has the advantage that several such trial runs may be performed over a relatively short period of time. In criminal cases, it is not unusual that different input data leading to correspondingly different probabilities of guilt, still each indicate "guilt beyond reasonable doubt".¹⁵

2 R v BLOM¹⁶

In this section, we apply the methods of Bayesian networks in an attempt to unravel a case that was heard almost eight decades ago.

This case was concerned with the death of a woman found shortly after 10 pm on 29 April 1938 beside the railway line some 24 kilometres from Graaff-Reinet after a train travelling from Rosmead to Graaff-Reinet ran over

¹⁵ Edwards 1991 13 *Cardozo LR* 1036–1055; Charniak "Bayesian Networks without Tears" 1991 12(4) *AI Magazine* 50 61–62.

¹⁶ *R v Blom* (1939) AD 188; see the quoted judgment for more details; Zeffert and Paizes *Essential Evidence* (2010) 27; Paizes "The Law of Evidence: Seven wishes for the next twenty years" 2014 27(3) *South African Journal of Criminal Justice* 272 272–292.

her head. Rosmead is a railway junction, situated 12 kilometres to the east of Middelburg in the Eastern Cape Province of South Africa. Besides the mutilation to her body caused by the train, there were no indications of previous physical injury. The victim was seen alive and well about an hour before the abovementioned incident. The *post-mortem* was conducted about 15 hours after her death by the district surgeon and a medical practitioner. Internally there was no trace of any poison. Because the head was mangled, it was not possible to determine the cause of death. She may have died earlier as a result of a stab or blow to the head after which the dead body was placed on the rail in such a way that the cause of death would not be detectable after the train ran over it. Earlier in the day, the victim died the defendant bought an ounce of chloroform in Graaff-Reinet, signing a false name in the poison register. In those days, chloroform was known to be used as an insecticide, and this was given as the reason for the purchase. She could have died by the application of chloroform to the face but all traces of such inhalation would have evaporated by the time the *post-mortem* took place. It was noted that the quantity of blood spilt was less than one might have expected if the victim was alive when she was placed on the railway line. This raised the possibility that she died as a consequence of the application of chloroform, or a stab or impact to the head.

At about 5.30 pm on that day, the defendant was said to have been seen cycling some distance to the north-east of Graaff-Reinet and 13 kilometres from the spot where the victim's body was found. However, the defendant's family members and other witnesses testified that on the evening and night in question, the defendant was at his brother's farm 3 kilometres east of Graaff-Reinet and he could therefore not have been at the spot where the victim met with her death. The trial court rejected the defendant's alibi.

The trial court stated that the victim's death was caused by the administration of chloroform "or some substance of similar properties". Clearly, the trial court did not arrive at an exact conclusion about the cause of the victim's death. The defendant (who was found guilty by the trial court) bought specifically chloroform. The Appeal Court overturned this and specified that chloroform and nothing else was the cause of death.

During the court proceedings, the defendant refrained from giving evidence, possibly on the advice of counsel. It is not known why – whether he felt uneasy in the courtroom surroundings, and could therefore, come across as being guilty or whether, he may have been protecting someone else, *et cetera*. It was also alleged that the defendant had a relationship with the victim but the court rejected this evidence on grounds of hearsay.

In short, there was no direct evidence linking the victim's death to the defendant. The case rested entirely on circumstantial evidence on grounds of which the defendant was found guilty of murder. It should be mentioned that at the time of the trial the death sentence was a possible consequence of judicial proceedings for murder.

Remark: We remind ourselves of our uneasy relationship with coincidence. Coincidence is a common occurrence in all walks of life. Often things just happen concurrently. It is quite possible that different pieces of evidence that seem to point in the same direction do so coincidentally, and for no reason

at all.¹⁷ But for some strange reason, many people believe that such occurrences are strong indications of guilt. Incorrect conclusions may be drawn, believing that events are somehow beyond coincidence. Such conclusions can only be drawn after careful analysis of the combination of probabilities of the respective pieces of evidence has been performed in conformance with the principles of probability theory. To illustrate we contemplate the innocuous pastime of coin flipping.

If an unbiased coin is flipped six times, then the probability that somewhere in the sequence of outcomes the same side of the coin will be observed in a run of five or more consecutive trials¹⁸ is 9.3%. Likewise, if an unbiased coin is flipped a hundred times, then the probability that somewhere in the sequence of outcomes the same side of the coin will be observed in a run of five or more consecutive trials is a remarkable 97.1%.¹⁹ This shows that clusters of coincidences are substantially more probable in larger populations.²⁰ Back to *R v Blom*.

Example 2

As far as can be gathered from the recorded judgment the possibility that some other person (other than the defendant) might have been the perpetrator was not given much consideration over the course of the judicial process. Watermeyer JA mentioned *en passant* in his contribution to the judgment the possibility that someone else might have been the perpetrator, but this line of thought was apparently not pursued in depth. It is respectfully submitted that it seems as if the collective feeling of the court might rather have been the following:

The chances of finding this evidence in an innocent man are so small that you can safely disregard the possibility that this man is innocent.²¹

However, this is indeed an error of logic, known as the *prosecutor's fallacy*.²²

Using official information for the year 1938²³ the adjacent districts of Middelburg and Graaff-Reinet could have had about 16,000 male inhabitants, excluding children. Reasons for the possibility that there might have been someone else that murdered the victim

“was the evidence that suggested that she had other lovers since she had two illegitimate children, as well as the evidence that she lived alongside a public

¹⁷ Muller 2014 35(2) *Obiter* 173 173–187.

¹⁸ The word *trial* is used as in probability theory. In this case, it refers to a single flip of the coin.

¹⁹ Berresford “Runs in Coin Tossing” 2002 33 *The College Math Journal* 391 391–393.

²⁰ Muller 2014 35(2) *Obiter* 173–187.

²¹ Fenton and Neil 2011 36 *Australian Journal of Legal Philosophy* 127.

²² The name derives from police and prosecutors that seek to find evidence that fits their theory as opposed to developing a theory based on existing evidence.

²³ Union Office of Census and Statistics “District Statistical Summary” 1938 19 *Official Year Book of the Union* viii.

road, and that she had met and talked to a man that night about an hour before her body was found on the line".²⁴

A reasonable estimate of the probability that a randomly chosen male person from the region satisfies the evidence associated with the perpetrator could be about 0.05%, that is, one in 2000. (This could also equivalently be formulated as follows: the probability that a randomly chosen male person from the region does *not* satisfy the evidence associated with the perpetrator could be about 99.95%, that is, 1999 in 2000.)

Let A and B be the following events:

A : The evidence associated with the perpetrator applies to the defendant.

B : The defendant is innocent.

Consider the following two conditional probabilities:

$P(A|B)$ = The probability of A , given B (that is, the probability that the evidence associated with the perpetrator applies to the defendant, given the fact that the defendant is innocent).

$P(B|A)$ = The probability of B , given A (that is, the probability that the defendant is innocent, given the fact that the evidence associated with the perpetrator applies to the defendant).

Failing to distinguish between $P(A|B)$ and $P(B|A)$ constitutes a serious error of logic in handling uncertainty which could have grave consequences for the defendant. Appendix (b) shows that two conditional probabilities with transposed conditionals may, in general, have widely different numerical values.

People easily come to the mistaken conclusion that the abovementioned probability of 0.05% is the probability that the defendant is innocent. This frequently happens when probabilities are considered, and it is also referred to as the *prosecutor's fallacy*.²⁵ Although intuitively appealing, it does not imply that the defendant is the perpetrator. Other evidence is required to prove that. The probability 0.05% refers to the conditional probability $P(A|B)$. We are more concerned with $P(B|A)$.

In the context of Example 2 we may understand the probability 0.05% as follows: given a male population of 16,000 in the region, there may have been up to $0.0005 \times 16,000 = 8$ individuals who could each have been the perpetrator. We have a cluster of 8 coincidences. The advantage of such information is that it narrows down the size of the pool of potential suspects, but usually we do not know who the other people in the pool are. However, the defendant will be a potential suspect. The court should decide on the basis of further evidence whether the defendant is indeed the guilty person among all the other potential suspects. So the probability that the defendant is the perpetrator is about $1/8 = 12.5\%$.

²⁴ Zeffertt and Paizes *Essential Evidence* 27.

²⁵ Aitken *Statistics and the Evaluation of Evidence for Forensic Scientists* (1995) 36–38; Muller 2012 23(3) *Stell LR* 599 600–601.

Example 3

Combining the above, we have the following Bayesian network:

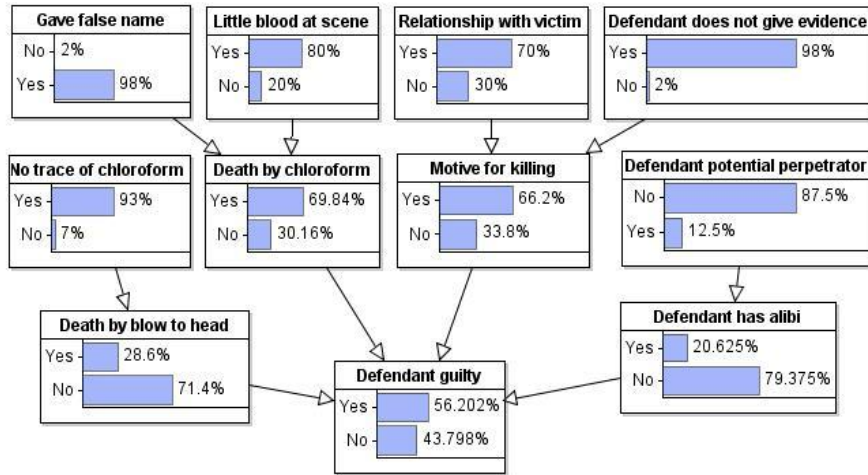


Figure 5

The parent nodes in figure 5 are “No trace of chloroform”, “Gave false name”, “Little blood at scene”, “Relationship with victim” and “Defendant does not give evidence”. The assigned node probability table for each of them appears in figure 5. The node “Defendant potential perpetrator” is also a parent node and it was obtained in Example 2. The node probability tables for the other nodes (underlined in each instance) are the following:

Gave false name:	Yes	No		
Little blood at scene:	Yes	No	Yes	No
<u>Death by chloroform:</u>				
Yes	0.7	0.3	0.8	0.3
No	0.3	0.7	0.2	0.7

Relationship with victim:	Yes	No		
Defendant does not give evidence:	Yes	No	Yes	No
<u>Motive for killing:</u>				
Yes	0.7	0.3	0.6	0.2
No	0.3	0.7	0.4	0.8

No trace of chloroform:	Yes	No		
<u>Death by blow to head:</u>				
Yes	0.3	0.1		
No	0.7	0.9		

Defendant potential perpetrator:	Yes	No		
<u>Defendant has alibi:</u>				
Yes			0.15	0.6
No			0.85	0.4

Figure 6

The last node probability table (figure 7) needed is “Defendant guilty”. This node table calculates the final probability of the defendant’s guilt:

Motive for killing:	Yes							
Death by chloroform:	Yes				No			
Death by blow to head:	Yes		No		Yes		No	
Defendant has alibi:	Yes	No	Yes	No	Yes	No	Yes	No
<u>Defendant guilty:</u>								
Yes	0.7	0.8	0.7	0.8	0.7	0.8	0.2	0.3
No	0.3	0.2	0.3	0.2	0.3	0.2	0.8	0.7

Motive for killing:	No							
Death by chloroform:	Yes				No			
Death by blow to head:	Yes		No		Yes		No	
Defendant has alibi:	Yes	No	Yes	No	Yes	No	Yes	No
<u>Defendant guilty:</u>								
Yes	0.2	0.4	0.3	0.5	0.17	0.37	0.03	0.1
No	0.8	0.6	0.7	0.5	0.83	0.63	0.97	0.9

Figure 7

Figure 5 shows that the defendant is guilty with estimated probability about 56.2%. This does not constitute “guilt beyond reasonable doubt”.

3 CONCLUSION

In section 2 we followed a typically Bayesian path. Initially (with meagre information) we estimated that the defendant was no more likely being guilty than any other male member of the population.

Then we updated this in Example 2 and found that the probability that the defendant was guilty was 12.5%.

In Example 3, we finally used the previously obtained probability as well as other information and obtained the posterior probability 56.2% that the defendant was guilty. This result is remarkably stable: if the values in the node probability tables were reasonably adjusted, then the posterior probability of guilt would indeed be different, but still far from “guilt beyond reasonable doubt”. The defendant should therefore not have been found guilty.

APPENDIX

The manual calculations (a), (b) and (c) mentioned in Example 1 are here performed from first principles.

- (a) We calculate $P(H)$, the probability that a randomly selected person from this group suffers heart failure.

$$\begin{aligned}
P(H) &= P(D \cap H) + P(\sim D \cap H) \\
&= P(D \cap S \cap H) + P(D \cap \sim S \cap H) + P(\sim D \cap S \cap H) + P(\sim D \cap \sim S \cap H) \\
&= P(D \cap S) P(H|D \cap S) + P(D \cap \sim S) P(H|D \cap \sim S) + \\
&\quad P(\sim D \cap S) P(H|\sim D \cap S) + P(\sim D \cap \sim S) P(H|\sim D \cap \sim S) \\
&= P(D) P(S) P(H|D \cap S) + P(D) P(\sim S) P(H|D \cap \sim S) + \\
&\quad P(\sim D) P(S) P(H|\sim D \cap S) + P(\sim D) P(\sim S) P(H|\sim D \cap \sim S) \\
&= (0.2)(0.1)(0.75) + (0.2)(0.9)(0.6) + \\
&\quad (0.8)(0.1)(0.55) + (0.8)(0.9)(0.3) \text{ according to figure 1} \\
&= 0.383 \text{ (= 38.3\%).}
\end{aligned}$$

- (b) We calculate $P(S|H)$, the conditional probability that a randomly selected person from this group who suffered heart failure is a smoker.

$$\begin{aligned}
P(S|H) &= \frac{P(S \cap H)}{P(H)} = \frac{P(S \cap H \cap D) + P(S \cap H \cap \sim D)}{P(H)} \\
&= \frac{(0.2)(0.1)(0.75) + (0.8)(0.1)(0.55)}{0.383} \text{ as in (a)} \\
&= 0.154 \text{ (= 15.4\%).}
\end{aligned}$$

(It can similarly be shown that $P(H|S) = 0.59 = 59\%$.)²⁶

- (c) We calculate $P(D \cup S)$, the probability that a randomly selected person from this group indulges in at least one of the following: an unhealthy diet (D), or smoking (S).

$$\begin{aligned}
P(D \cup S) &= P(D) + P(S) - P(D \cap S) \\
&= 0.2 + 0.1 - (0.2)(0.1) \text{ since } D \text{ and } S \text{ are mutually} \\
&\quad \text{independent events} \\
&= 0.28 \text{ (= 28\%).}
\end{aligned}$$

²⁶ Note that $P(S|H)$ and $P(H|S)$ have different numerical values.