Ian W. Evett,[1] *D.Sc.; Peter D. Gill,*[2] *Ph.D.; Graham Jackson,*[3] *B.Sc.; Jonathan Whitaker,*[2] *Ph.D.; and Christophe Champod,*[2] *Ph.D.*

Interpreting Small Quantities of DNA: the Hierarchy of Propositions and the Use of Bayesian Networks

REFERENCE: Evett IW, Gill PD, Jackson G, Whitaker J, Champod C. Interpreting small quantities of DNA: the hierarchy of propositions and the use of Bayesian networks. J Forensic Sci 2002;47(3):520–530.

ABSTRACT: The dramatic increase in the sensitivity of DNA profiling systems that has occurred over recent years has led to the need to address a wider range of interpretational problems in forensic science. The issues surrounding questions of the kind "whose DNA is this?" have been the subject of considerable controversy but now it is clear that the emphasis is shifting to questions of the kind "how did this DNA get here?" Such issues are discussed in this paper and new insights are provided by two particular recent developments. First, the notion of the "hierarchy of propositions" that has arisen from a project called Case Assessment and Interpretation (CAI) that has been running in the British Forensic Science Service (FSS). Second, a technique for drawing inferences in the face of many interacting considerations, known as "Bayesian networks"—or "Bayes' nets" for short—that has been the subject of an earlier paper in this journal (1). The discussion is carried out by means of case studies, based on actual cases. It is clear that, whereas the inference in relation to the source of the DNA in a crime sample might be overwhelmingly strong, the inference in relation to the propositions that a jury must consider relating to the identity of the actual offender may be much more tentative.

KEYWORDS: forensic science, interpretation, pre-assessment, likelihood ratios, Bayesian networks, DNA profiling

Much of the DNA debate of the last decade has been concerned with the process of individualization, where the considerations of statistics and population genetics are paramount. Now it is apparent that the emphasis in court will shift from questions of the kind "whose DNA is this?" to "how did this person's DNA get here?". Such questions invoke the standard forensic considerations of recovery, transference, persistence and contamination; however, the ways in which these factors will be viewed will develop as it becomes possible to generate profiles from ever-smaller quantities of DNA. Already it is feasible to amplify the DNA molecule from a single cell.

In this paper we show how the inferential issues that follow from these exciting developments may be clarified by means of two particular aids to logical reasoning. The first of these, which we refer to as *the hierarchy of propositions,* arose from a Case Assessment and Interpretation (CAI) project that has been running for three years within the Forensic Science Service (FSS). The work of the project has resulted in a series of papers (2–5) and the concept of the hierarchy of propositions is discussed in some detail in the second and fourth of these: we outline the idea in the next section. The second aid, Bayesian networks, which we will discuss, has been described in a previous paper in this journal (1) and in this paper we show the application of Bayesian networks software to two case studies.

The Hierarchy of Propositions

The hierarchy embodies the principle that interpretation of evidence is not possible unless one considers at least two competing propositions. In most cases these will represent the prosecution and defense positions. At court, the jury will consider propositions of the kind:

- The defendant raped the victim
- The victim was raped by some unknown person

For obvious reasons, we refer to these as *Offence level* propositions: this is the third, or highest level in the hierarchy. Propositions at this level generally invoke considerations quite outside the competence of the expert witness and very often it will be necessary for the scientist to address propositions at the second or *Activity level*, for example

- The defendant is the person who smashed the window
- The defendant has never been at the scene

For a scientist to address activity level propositions, it will be necessary for him/her to have a body of information relating to the alleged incident and whatever the suspect/defendant is saying (if anything): we call this the *framework of circumstances*. In the present instance, the first proposition is conditioned by the information that the window was broken by a man and the second proposition is here conditioned by what the defendant is saying—that he has never been at the scene. The propositions are liable to change if the framework changes in any way; for example, if the defendant later admits being at the scene but claims that he was an innocent bystander, then the second proposition would change.

The framework of circumstances also conditions the interpretation of the evidence in relation to the two propositions. For exam-

[1] The Forensic Science Service, 109 Lambeth Road, London SE1 7LP.
[2] The Forensic Science Service, Trident Court, Solihull Parkway, Birmingham Business Park, Solihull B37 7YN.
[3] The Forensic Science Service, Washington Hall, Euxton, Chorley, PR7 6HJ.

Received 21 Feb. 2001; and in revised form14 Nov. 2001; accepted 14 Nov. 2001.

ple, in the present instance the quantity of matching glass that might be found on the clothing of the defendant would depend on factors such as: the size and location of the window; how the window was allegedly smashed; and the length of time between the incident and the taking of the clothing for examination. So it is when the scientist is addressing propositions at this level that it is necessary for him/her to consider issues of transference and persistence. If, however, the framework contains so little information that the addressing of these issues becomes problematic then the scientist will need to move to what we call Level 1 or *Source level* propositions, such as:

- The glass recovered from the defendant's clothing came from the broken window
- The glass recovered from the defendant's clothing came from some other source.

At this level, the problem of addressing transfer and persistence must be left to the court where, no doubt, advocacy will play a major role.

There is another consideration that had not been taken into account when the original hierarchy (2) was drawn up. Classically, at the source level, the scientist might have addressed propositions of the kind:

- The crime stain was left by the suspect
- The crime stain was left by some unknown person

This pair would be conditioned by the scientist's confidence that the genotyping information had actually been derived from the stain that was visible. However, given the sensitivity of today's technology it is possible to envisage cases in which such an inference is dubious. The visible stain, for example, might be very small and degraded, yet the substrate bears another, invisible deposit from another source that is of better condition and readily amplifiable. Given such conditions, it may be necessary to move to what we now call *sub-Level 1 propositions,* such as:

- The DNA is that of the suspect
- The DNA is that of some unknown person

As a general rule, it is felt that the higher the level of the propositions that the scientist can address, the greater the value that is added to the criminal justice process. When the scientist has to settle for the lower levels then the task of moving to the high-level offence propositions must be left to the jury.

Case Study 1

A family returned to their home to find that it had been burgled during their absence. It appears that the intruder (or intruders) gained entry by means of breaking a window at the rear of the property. Near to the point of entry, on a paved patio, a crime scene officer recovered a cigarette end. None of the family smoked, the surroundings of the house were carefully maintained and the patio regularly swept. A few weeks later, a suspect was apprehended in connection with another incident and he was found to have in his possession a checkbook for drawing on the account of one of the family members. He said that he received the checkbook from "an unknown man in the pub." He denied ever being anywhere near to the home that was burgled.

If the cigarette end is submitted for DNA profiling and the results compared with the suspect's DNA then various outcomes are possible, which we categorise as follows:

Match—a single profile from the cigarette end that is attributable to the same genotype as that of the suspect.

Mixture/Match—a mixed profile from the cigarette end that includes alleles that are present in the profile of the suspect.

Different—This includes either a single profile that is a different genotype from that of the suspect; or a two person mixture that differs from the suspect's DNA profile.

No profile—Our current technique fails to detect any DNA profile.

The CAI project places considerable emphasis on the notion of *pre-assessment* (see, in particular, the first paper in the series (3))—that is, thinking in careful detail of the expected outcomes of an analysis beforehand. In the present case, referring to the suspect as X, we formulate the following pair of provisional activity level propositions:

$\underline{S}$: X is the person who smoked the cigarette

$\overline{S}$: Some other unknown person smoked the cigarette

In the event of a match, we might consider sub-level 1 propositions, of the kind:

The DNA recovered from the cigarette end came from the suspect

The DNA recovered from the cigarette end came from some unknown person

The likelihood ratio for this pair of propositions would simply be the inverse of the match probability. For the purpose of this discussion, we will take the match probability to be essentially zero, so the infinite likelihood ratio would correspond to the kind of categorical opinion that is usually associated with fingerprints opinions: the DNA is that of the suspect. Of course, this is not the way that DNA evidence is interpreted in practice, but it simplifies the discussion without detracting from the principles.

We will now expand the analysis to take account of three issues that we present in the form of paired propositions:

$\underline{A}$: DNA from X entered the process by innocent means

$\overline{A}$: No DNA from X entered the process by innocent means

Such adventitious transfer may flow from various causes, including contamination as part of the analytical procedures.

We introduce the consideration that is equivalent to the *relevance* term as used by Stoney (6) and Evett (7):

B: The person who smoked the cigarette left sufficient DNA to give a profile

$\overline{B}$: The person who smoked the cigarette did not leave sufficient DNA to give a profile

We also consider:

$\underline{P}$: DNA from some third person entered the process

$\overline{P}$: No DNA from a third person has entered the process

Here we use "third person" to convey the idea of someone other than the suspect and also someone other than the person who smoked the cigarette, if that is not the suspect. For the purpose of this discussion, we will assume that the uncertainties of each of the three pairs of propositions just defined are not influenced by each other. For example, whether or not DNA from X entered the process by innocent means is not influenced by whether or not the person who left the cigarette left sufficient DNA to give a profile; nor by whether or not DNA from some third person entered the pro-

cess. These independence assumptions are not essential, but they do make the discussion much simpler.

Let E denote the outcome of the comparison, then we seek the likelihood ratio:

$$LR = \frac{\Pr(E \mid S, I)}{\Pr(E \mid \overline{S}, I)}$$

It is now necessary to expand the terms in the numerator and denominator to take account of the considerations embodied in the three intermediate pairs of propositions. This can be done by invoking the "law of total probability" (explained, for example, in (8)). The existence of three pairs of propositions means that there are eight combinations of possibilities, such as, for example: *A, B,* and *P* are all true; *A* is false, B and P are both false; and so on. If we take the conditioning on I to be implicit, then the numerator is then written out as:

$$\begin{aligned}&\Pr(E \mid ABPS)\Pr(ABP \mid S) + \Pr(E \mid \overline{A}BPS)\Pr(\overline{A}BP \mid S)\\ &+ \Pr(E \mid A\overline{B}PS)\Pr(A\overline{B}P \mid S) + \Pr(E \mid \overline{AB}PS)\Pr(\overline{AB}P \mid S)\\ &+ \Pr(E \mid AB\overline{P}S)\Pr(AB\overline{P} \mid S) + \Pr(E \mid \overline{A}B\overline{P}S)\Pr(\overline{A}B\overline{P} \mid S)\\ &+ \Pr(E \mid A\overline{BP}S)\Pr(A\overline{BP} \mid S) + \Pr(E \mid \overline{ABP}S)\Pr(\overline{ABP} \mid S)\end{aligned}$$

TABLE 1—(a) *Pr(E|...S) terms and* (b) *Pr(E|...$\overline{S}$) terms.*

	Match	Mixture/ Match	Different	No Profile
Pr(E| $ABPS$)	0	1	0	0
Pr(E| $\overline{A}BPS$)	0	1	0	0
Pr(E| $AB\overline{P}S$)	1	0	0	0
Pr(E| $\overline{A}B\overline{P}S$)	1	0	0	0
Pr(E| $A\overline{B}PS$)	0	1	0	0
Pr(E| $\overline{A}\overline{B}PS$)	0	0	1	0
Pr(E| $A\overline{BP}S$)	1	0	0	0
Pr(E| $\overline{ABP}S$)	0	0	0	1

	Match	Mixture/ Match	Different	No Profile
Pr(E| $ABP\overline{S}$)	0	1	0	0
Pr(E| $\overline{A}BP\overline{S}$)	0	0	1	0
Pr(E| $AB\overline{P}\overline{S}$)	0	1	0	0
Pr(E| $\overline{A}B\overline{P}\overline{S}$)	0	0	1	0
Pr(E| $A\overline{B}P\overline{S}$)	0	1	0	0
Pr(E| $\overline{A}\overline{B}P\overline{S}$)	0	0	1	0
Pr(E| $A\overline{BP}\overline{S}$)	1	0	0	0

TABLE 2—*Pr(ABP| .) terms.*

Pr(ABP| .)	bpa
Pr($\overline{A}BP$| .)	$bp(1-a)$
Pr($AB\overline{P}$| .)	$b(1-p)a$
Pr($\overline{A}B\overline{P}$| .)	$b(1-p)(1-a)$
Pr($A\overline{B}P$| .)	$(1-b)pa$
Pr($\overline{A}\overline{B}P$| .)	$(1-b)p(1-a)$
Pr($A\overline{BP}$| .)	$(1-b)(1-p)a$
Pr($\overline{ABP}$| .)	$(1-b)(1-p)(1-a)$

The eight terms of the kind $\Pr(E \mid \ldots.)$ can be simply assigned as is shown in Table 1(a). For example, if A, B, and P are all true, then we will expect to see DNA from the suspect, by both innocent and offence related means, as well as DNA from a third person: the observation will thus be a mixture/match with probably one. Readers may find it interesting and instructive to check the other lines of this table.

We now define the following terms:

a is the probability that DNA entered the process by innocent means and we assume that this occurrence is independent of whether or not S is true:

$$\Pr(A \mid S) = \Pr(A \mid \overline{S}) = \Pr(A) = a$$

b is the probability that the person who smoked the cigarette left sufficient DNA to give a profile. Without any particular knowledge of whether or not the suspect is a good shedder of DNA, this probability is the same whether or not S is true:

$$\Pr(B \mid S) = \Pr(B \mid \overline{S}) = \Pr(B) = b$$

p is the probability that DNA from some third person entered the process. There is no reason to believe that this is dependent on the truth or otherwise of S.

$$\Pr(P \mid S) = \Pr(P \mid \overline{S}) = \Pr(P) = p$$

Then, bearing in mind that we assume the three pairs of intermediate propositions to be independent, the terms of the kind $\Pr(ABP|S)$ can be written out as in Table 2.

For the denominator, the eight $\Pr(E \mid \ldots.)$ terms are shown in Table 1(b). The terms of the kind $\Pr(ABP \mid \overline{S})$ are the same as the corresponding terms in Table 2.

It is now possible to consider the pre-assessment of the case. We have allowed for four possible outcomes, E, and Tables 1 and 2 enable us simply to work out the probability of each outcome given that either S or $\overline{S}$ is true: these are shown in the second and third columns of Table 3 and follow from the appropriate cells of Tables 1 and 2 by means of the formula derived using the law of total probability. For example, if E is a "Match" then the numerator, $\Pr(E \mid \overline{S})$, is given by the addition of the probabilities on the 3rd, 4th, and 7th lines of Table 2 because they are the lines of Table 1(a) that are one

TABLE 3—*Pre-assessment table.*

Outcome, E	$\Pr(E \mid S)$	$\Pr(E \mid \overline{S})$	LR
Match	$(1-p)\{b + a(1-b)\}$	$(1-b)(1-p)a$	$1 + \frac{b}{a(1-b)}$
Mixture/match	$p\{b + a(1-b)\}$	$a\{b + p(1-b)\}$	$\frac{p\{b + a(1-b)\}}{a\{b + p(1-b)\}}$
Different	$(1-b)p(1-a)$	$(1-a)\{b + \mathrm{p}(1-\mathrm{b})\}$	$\frac{p(1-b)}{\{b + p(1-b)\}}$
No profile	$(1-b)(1-p)(1-a)$	$(1-b)(1-p)(1-a)$	1

and not zero i.e.:

$$\Pr(E = Match \mid S) = ab(1-p) + (1-a)b(1-p) + a(1-b)(1-p) = (1-p)\{b + a(1-b)\}$$

which is the term in the first row and first column of Table 3. All of the other terms in the second and third columns are similarly derived.

If we now take each term in the second column and divide it by the corresponding term in the third column, this gives the *LR* that we would calculate for the outcome corresponding to that row of the table. This is the pre-assessment table.

Clearly, our pre-assessment depends on the magnitudes we assign to the probabilities *a, b,* and *p*. If these values can be assigned, the pre-assessment can be carried out. For example, imagine that we believe that: the probability that the suspect's DNA entered the process for innocent means is small—say 0.01; the probability that the person who smoked the cigarette left DNA is moderate—say 0.9; and the probability that DNA from some third person entered the process is fairly small—say 0.1. Then the first line of Table 3 tells us that, if the suspect is truly the person who smoked the cigarette, then the probability that the outcome of the analysis will be a match is $0.9 \times 0.9 + 0.01 \times 0.1 = 0.81$ and the *LR* given that outcome would be approximately 900. This would denote the strength of the evidence supporting the proposition *S* that the suspect is the person who smoked the cigarette. Note that, although this is strong evidence, it is not in the orders of magnitude that courts have come to expect in DNA cases and this, as we have said, is because we are considering higher level propositions that relate to the activities that are of interest to a court. Recall that we have assumed an infinite *LR* for the sub-level 1 propositions relating to the origin of the DNA.

But, of course, we recognize that, in practice, it will be difficult to assign precise values to *a, b,* and *p*. They could, to a considerable extent, be informed by experimentation but will inevitably be case specific. The strategy then is to consider the *sensitivity* of the outcomes to ranges of values for the probabilities. We could, using Table 3, produce families of graphs but it is more informative to employ computer software to carry out simulations. Indeed, in more complex cases this is the only realistic way to proceed. This brings us to the discussion of Bayesian networks.

Use of Bayesian Networks

A Bayesian network is a graphical model for expressing the probabilistic relationships among a set of variables. In essence, a network consists of *nodes*, representing variables, and arrows that symbolize the conditional relationships between the variables. A variable might be a continuous variable, such as a measurement on a quantity; or a set of discrete outcomes such as the truth or otherwise of a set of propositions. The model efficiently encodes the joint probability distribution for a large inter-related set of variables. It gives not only a graphical representation of the problem but allows probabilistic inference.

Figure 1 shows the Bayesian network for our first case study that has been implemented in a program called Hugin Lite (Version 5.3, free demonstration version available at http:\\www.hugin.com). There are five nodes. Nodes 1 to 4 correspond to the four pairs of propositions that we have defined in the previous section; each of these has two states, corresponding to which of the alternative propositions is true. Node 5, which has four possible states, corresponds to the outcome of the analysis and its dependence on Nodes 1 to 4 is symbolised by the four arrows. Note that Nodes 1 to 4 have no arrows between them, because of the independence assumptions that we have made but the network could easily be modified to allow for dependencies. For example, if it were known that the suspect was a particularly heavy shedder of DNA then an arrow from Node 1 to Node 2 would enable us to include different values for the probability of the smoker leaving DNA, depending on whether *S* or $\bar{S}$ were the case. The computer program makes such modifications simple, whereas the algebraic solution soon becomes extremely complex.

Once the network has been drawn, and the appropriate conditional probabilities specified, then it can be compiled, just like any other high level programming language. Figure 2 shows the result of compiling the program with the values of $a = 0.01$, $b = 0.9$, and $p = 0.1$ that were used in the previous section. For the time being, the probabilities for *S* and $\bar{S}$ have been set to 0.5—this will be changed shortly. The main window shows the network and the left hand window shows the probabilities, expressed as percentages.

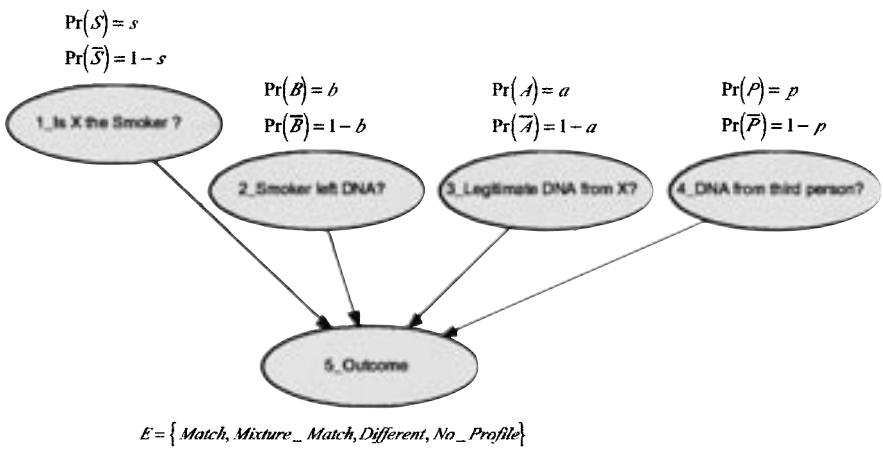

FIG. 1—*Bayesian network for the Case Study 1.*

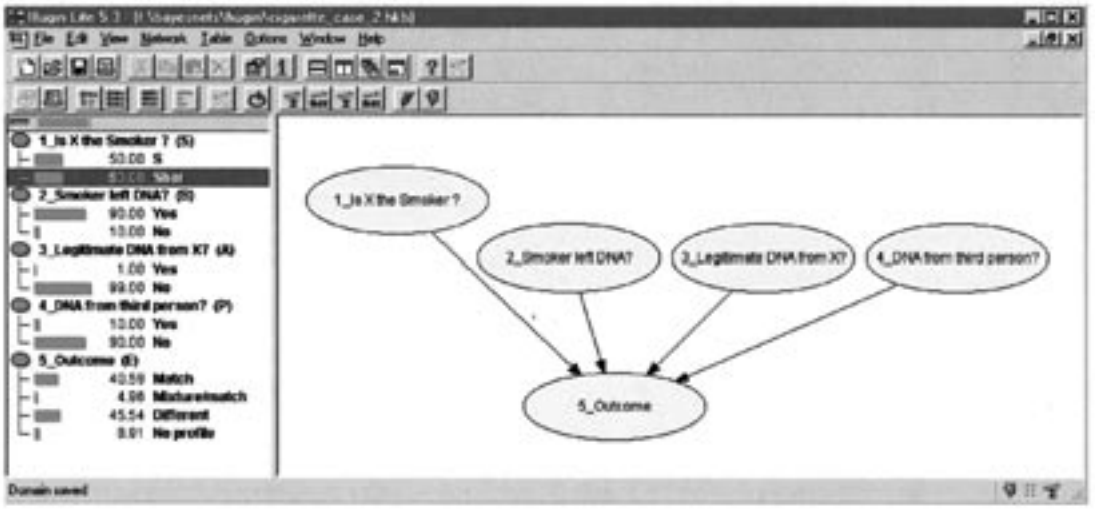

FIG. 2—*Bayesian network for Case Study 1 in Hugin.*

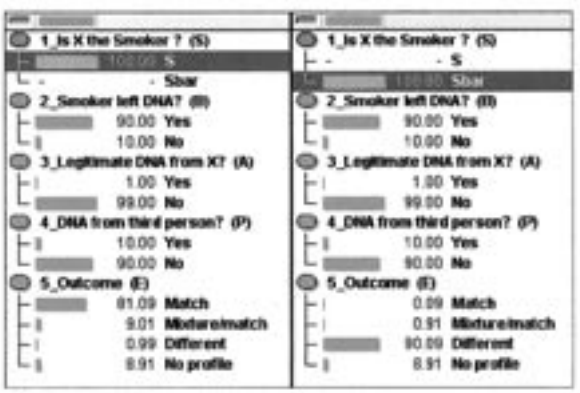

FIG. 3—*Pre-assessment in Hugin for Case Study 1.*

The power of the Bayesian network is best realized when it is used for "back propagation": that is, given a particular outcome, such as a match, the program can be used to update the probabilities for S and $\bar{S}$. However, at present, we are emphasising the pre-assessment process and this may be done as follows. First, we set the probabilities at Node 1 to 100% and zero for S and $\bar{S}$ respectively then, when the program is run, the probabilities for the four outcomes given that S is true are calculated. These can be seen in the left hand column of Fig. 3 and correspond to the second column of Table 3, calculated with the assigned values of a, b, and p. Note that the probability of the outcome "match" is 0.81, as calculated earlier. Next, if the probability $\bar{S}$ is set to 100% and the table recalculated, the second column of Table 3 is derived, showing the probabilities of the four outcomes given that $\bar{S}$ is true.

The *LR's* for the four outcomes can be calculated by taking the ratios of the corresponding probabilities in the four rows.

Recall that the intention is to carry out a sensitivity analysis, however. The scientist may be reluctant to provide precise values for the probabilities and it is useful now to regard a, b, and p as parameters whose values are unknown. The scientist then expresses his/her knowledge in the form of probability intervals for the three parameters. Imagine that the scientist expresses his views in the form of 80% probability intervals for a, b, and p as in Table 4.

TABLE 4—*Lower bound and upper bound for the parameters.*

Parameter	Lower Bound	Upper Bound
b	0.8	0.95
a	0.001	0.01
p	0.05	0.1

These beliefs can usefully be modelled by *Beta distributions* as shown in Fig. 4.

Our aim is now to build the uncertainties in a, b, and p into our pre-assessment. Whereas Hugin Lite could be used manually to explore the sensitivity of the outcome to changing values of a, b, and p, this would be a cumbersome process. It is far more effective to carry out simulations and this is not possible within this version of the program. Consequently, we have implemented the network in MATLAB 5.0 (Mathworks, Inc.) using the Bayes Net Toolbox 2.0 developed by K. Murphy (http://www.cs.berkeley.edu/~murphyk/). This implementation leads to the same output as Hugin but it is more versatile.

Using this program, it is possible to carry out large number of pre-assessments, using values of a, b, and p generated at random from the Beta probability distributions shown in Fig. 4. This is done by means of an appropriately programmed random number generator: for each simulation, a value of a is generated from the probability distribution in Fig. 4(a) and by analogy for b and p. One thousand simulated pre-assessments were carried out in this way, computing at each iteration the probabilities for each outcome given S and $\bar{S}$. The results are two probability distributions for the *LR*, one given that the prosecution proposition is true (S), the other given that the defense proposition is true ($\bar{S}$). These distributions are shown in Fig. 5 and summarized in Table 5.

Consider the first graph in Fig. 5. Recall that we are at the pre-assessment stage so we do not yet know the outcome of the analysis so we are not yet in a position to calculate the *LR*—we can only talk about a probability distribution for the *LR*. This first graph is a probability distribution for the log (base 10) of the *LR* given that the prosecution proposition (S) is actually true. From the height of the

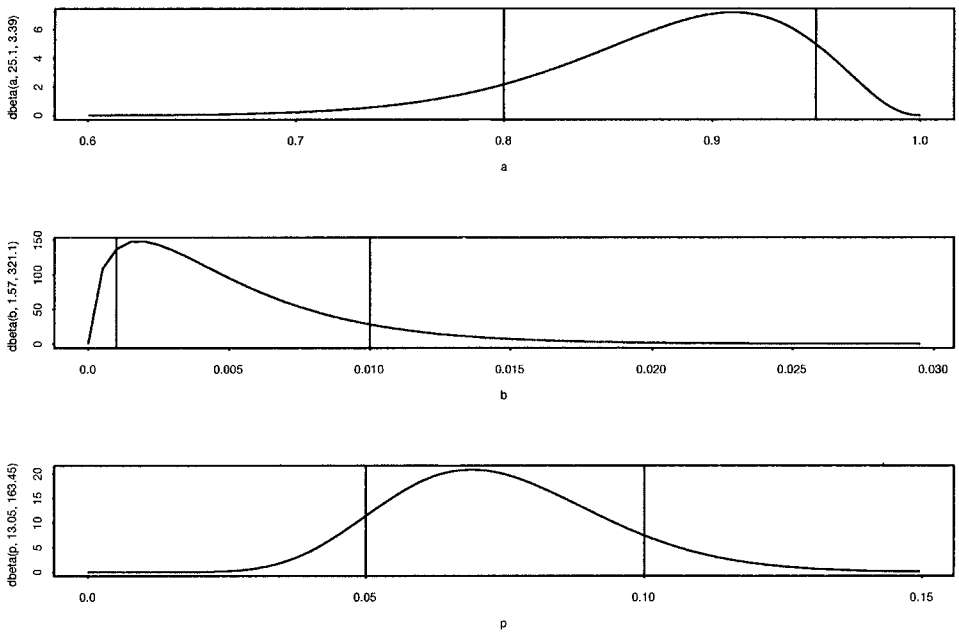

FIG. 4—*Beta distributions modelling the uncertainty on parameters a, b, and p. The vertical bars show the upper and lower bounds.*

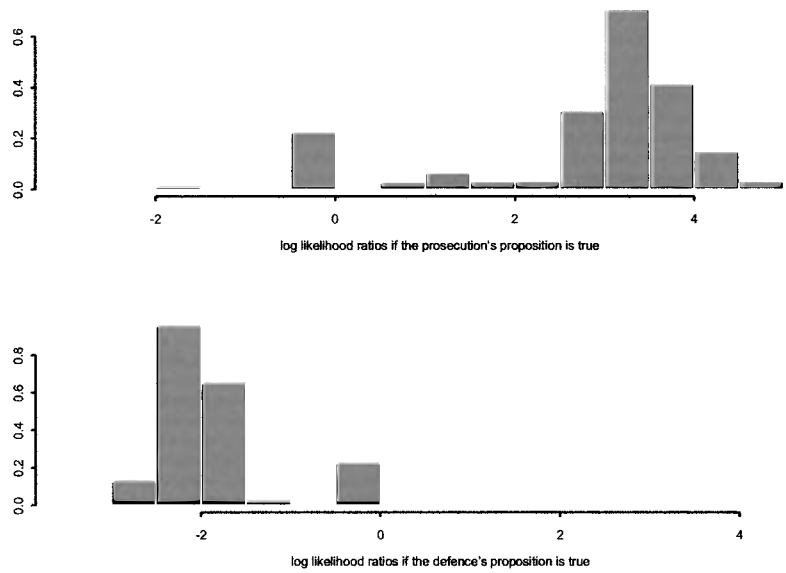

FIG. 5—*Distribution of likelihood ratios in Case Study 1 following 1000 pre-assessment simulations.*

TABLE 5—*Summary of the pre-assessment in Case Study 1.*

Likelihood Ratios	Verbal Equivalent	If *S* Is True (%)	If $\overline{S}$ Is True (%)
< 0.001	Strong evidence to support $\overline{S}$	0	54.6
0.001 to 0.01	Moderately strong evidence to support $\overline{S}$	0.35	33.7
0.01 to 0.1	Moderate evidence to support $\overline{S}$	0.55	11.2
0.1 to 1	Limited evidence to support $\overline{S}$	11.2	0.22
1 to 10	Limited evidence to support S	1.3	0.22
10 to 100	Moderate evidence to support S	4.6	0.03
100 to 1000	Moderately strong evidence to support S	16.8	0.03
> 1000	Strong evidence to support S	65.2	0

bars, it can be seen that there is, approximately, an 80% chance of the outcome being such that an *LR* in excess of 100 is derived. Such *LR's* would, in the Forensic Science Service, be expressed as moderately strong, or strong, support for *S*: the equivalence is shown in the third column of Table 5. We must note that there is also a chance (approximately 11%) that the outcome of the analysis will be such that it will support the alternative $\bar{S}$—even though *S* is the true state.

The second graph in Fig. 5 shows the probability distribution for the *LR* given that the defense proposition ($\bar{S}$) is actually true. With a probability of approximately 99.5% we would expect to see an outcome that supported that proposition. We must recognize that there *is* a chance, albeit small—0.5%—that the outcome would yield support for the prosecution proposition. Again, this can be seen in more detail in the fourth column of Table 5. The idea of pre-assessment is that such graphs and tables should be used to inform a discussion with the customer who is to pay for the analysis and to form a clear idea of the customer's requirement and expectations.

Now we move on from pre-assessment stage. Imagine that the analysis has been carried out and an unmixed profile has been obtained that matches the profile of the suspect. It is necessary now to ask the network to provide a different kind of analysis—given that we have observed a match, what is the *LR* in relation to the two propositions *S* and $\bar{S}$? But, of course, there is not a precise single value for the *LR* because of the uncertainties in *a*, *b,* and *p*. The same sort of simulation process as we have just described leads to a probability distribution for the *LR* given that we know the outcome. This is shown in the first part of Fig. 6: note that the modal value represents a *LR* of approximately 1800.

If, on the other hand, the result of the analysis is that a profile is recovered from the cigarette that is different from that of the suspect then a similar process of simulation analysis can be carried. This leads to the probability distribution for the *LR* that is shown in the second part of Fig. 6.

Whereas the distributions in Fig. 5 would be used for pre-assessment, the distribution here would be used to inform discussions between the scientist and counsel—whether prosecution or defense (preferably both)—before a trial. It could also enable the scientist in assisting a court in understanding the importance of the relevant issues. Once again, we emphasise the difference between the addressing of sub-level 1 and Level 2 propositions. In the event of a match, there would be an infinite *LR* in favour of the sub-level 1 proposition that the DNA recovered from the cigarette is that of the suspect. However, in relation to the *activity level* propositions *S* and $\bar{S}$ the *LR* is of the order 1000. It seems to us that it is essential for forensic scientists to recognize this distinction and to have the ability to explain the issues to a court of law.

Case Study 2

The second case study is more complex than the first and enables us to explain the ease with which increasing complexity is accommodated within a Bayesian network. It is extremely difficult to envisage a full algebraic solution to this case.

The circumstances are as follows: a watch with a broken strap was recovered at the scene of a rape and it was suspected that the watch had been worn by the offender. The watch was submitted for DNA profiling using LCN with a view to providing intelligence information regarding the possible wearer. In such an intelligence mode, it is more difficult to proceed to a pre-assessment because there is no simple pair of propositions. However, given the existence of the National DNA database it appeared a reasonable strategy to proceed with the analysis of three samples taken from three different areas (Areas 1 to 3) of the watchstrap. Area 1 was the outside surface of the strap at the point of damage, Areas 2 and 3 were two separate zones of the inside surface of the strap. Mixed DNA profiles were obtained from each of the areas examined:

Area 1: A major component matching the victim's profile and a minor male component.

Area 2: A major male component (corresponding to the minor component in Area 1) and a minor component matching the victim's profile.

Area 3: A major male component as in Area 2 with an additional minor component that did not correspond to the victim's or the suspect's profile.

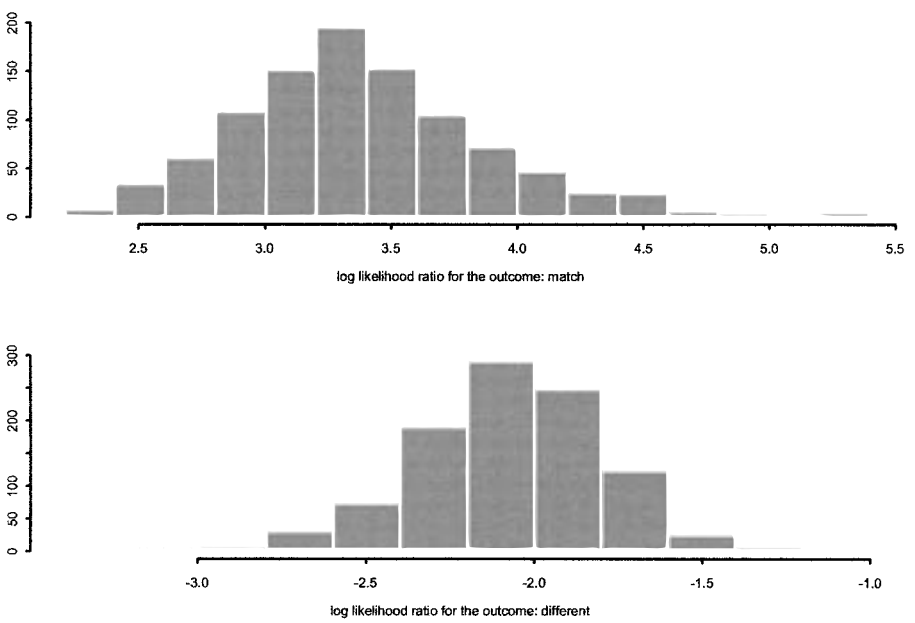

FIG. 6—*Distribution of likelihood ratios in Case Study 1 for the outcomes "match" and "different."*

The major male component was searched against the National DNA Database and a corresponding profile was found from a man who lived in the locality of the incident. This man became the main suspect. He admitted owning a watch, without specifically designating the recovered watch, but stated that he lost it approximately two to three weeks before the offence.

For sake of simplicity, we will treat the three areas from the watchstrap as one. Moreover, we will continue to assume that the DNA profiling has ultimate discriminating power: a match implies an individualisation, or an infinite *LR* at sub-Level 1. So the DNA component matching the victim's profile is identified and explained by the incident. The victim aside, the outcome is that the major component came from the suspect and the minor component came from an unknown person.

We will now address the impact of this result on the following pair of Level 3 propositions (offence level):

$\underline{O}$: X is the offender

$\overline{O}$: Some other unknown person is the offender

We know the outcome in this case and for the purpose of constructing the Bayes' network, we will focus on that outcome versus all other possible outcomes:

Observed outcome—a mixed profile from the strap that includes a major component matching the suspect's DNA profile and a minor component from one unknown person. This is the outcome of interest.

Any other outcome—This includes: a single matching profile; other mixture matches; different profiles; or no profile at all.

The breakdown of the list of outcomes could be more detailed, taking into account every possibility offered by the major and minor components. This would be recommended for carrying out a detailed pre-assessment stage. However, in our case, we would like to focus on one particular outcome and to assess the sensitivity of the Level 3 *LR,* given that outcome, to the various probabilities that are relevant to its assessment. To do this we first implemented a network in Hugin, shown in Fig. 7 that appeared to us to encapsulate the various intermediate considerations, and the dependencies that bore on the observed outcome of the comparison (Node 12).

The definitions of the nodes and their states are given in Table 6.

The conditional probability specifications are detailed and explained when necessary in Table 7.

TABLE 6—*Nodes and possible states of the Bayesian network in Case Study 2.*

Node Number	Node Description	Possible States
1	Is X the offender?	$\underline{O}$: X is the offender $\overline{O}$: X is not the offender
2	Did the offender wear the watch?	R: the watch was worn by the offender $\overline{R}$: the watch was not worn by the offender
3	If X is not the offender, has he ever worn the watch?	$\underline{W}$: X has worn the watch $\overline{W}$: X has never worn the watch
4	If X is not the offender, but has worn the watch, when was the last time?	T_1: the day of the offence T_2: a week ago T_3: a month ago T_4: never
5	Is there a possibility of scene or laboratory contamination?	$\underline{C}$: yes $\overline{C}$: no
6	Is there any contribution from a third person? By contribution from a third person, we mean all DNA contribution coming from sources other than the victim, X or the offender.	Th: yes $\overline{Th}$: no
7	Legitimate DNA profile from X	x_major; x_minor; no profile
8	Adventitious DNA profile from X	x_major; x_minor; no profile
9	DNA profile from a third person	z_major; z_minor; no profile
10	DNA profile from the offender	o_major; o_minor; x_major; x_minor; no profile
11	DNA profile from X unrelated with the offence	x_major; x_minor; no profile
12	Outcome from the DNA analysis of the crime sample following a comparison with the DNA profile from X	Match_Major/ Mix1_Minor Any other outcome

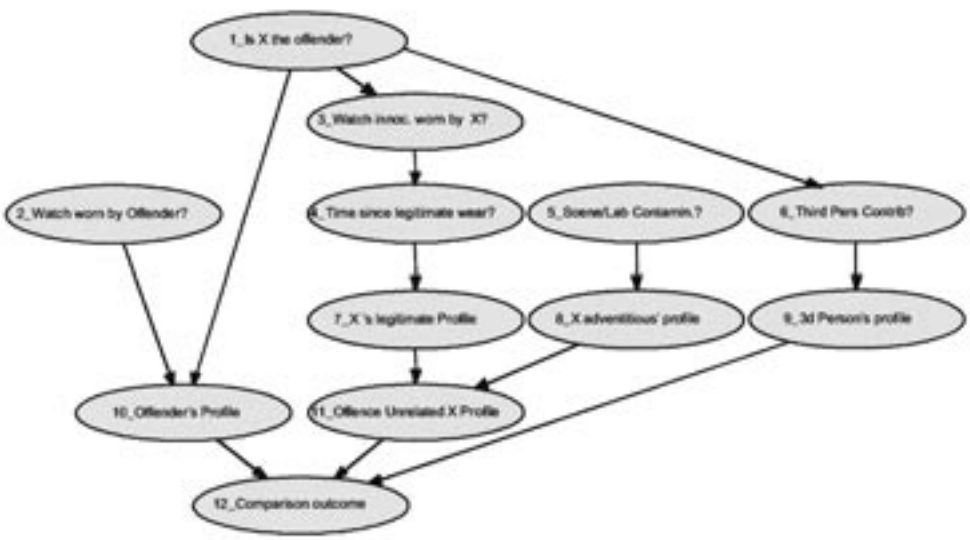

FIG. 7—*Bayesian network for Case Study 2.*

TABLE 7—*Conditional probability specifications for Case Study 2. When not stated explicitly, the probability for the complement of an event (with associated probability p) is simply 1-p.*

Probabilities					*Comments*
$\Pr(O) = 0.5$					The Bayesian network is initiated with this default value.
$\Pr(R) = r$					r represents our assessment of the relevance of the watch.
$\Pr(W\mid O) = 0$					The first probability is 0 by the definition of Node 3. The suspect admitted having a watch but there is doubt on whether it was that watch or another one. Evidence has been provided regarding paint splashes visible on the recovered watch and on a photograph showing a watch worn by the suspect. s represents our assessment of that evidence.
$\Pr(W\mid \overline{O}) = s$					
$\Pr(C) = c$					c represents our assessment of the probability of a contamination.
$\Pr(Th\mid O) = q$					q expresses our assessment of the probability that a third person contribution may have entered the process if X is not the offender. t is the probability for the same event but if X is not the offender.
$\Pr(Th\mid \overline{O}) = t$					
	W	$\overline{W}$			The probabilities (under W) were assessed based on the testimony of the suspect who admitted having lost his watch 2 to 3 weeks ago.
$\Pr(T_1\mid .)$	0	0			
$\Pr(T_2\mid .)$	0.5	0			
$\Pr(T_3\mid .)$	0.5	0			
$\Pr(T_4\mid .)$	0	1			
	T_1	T_2	T_3	T_4	These data represent our state of knowledge of the efficiency in function of the delay between the last wear and the analysis of our DNA profiling techniques on watchstraps.
$\Pr(x\_major\mid .)$	0.9	0.5	0.1	0	
$\Pr(x\_minor\mid .)$	0.05	0.3	0.2	0	
$\Pr(no\ profile\mid .)$	0.05	0.2	0.7	1	
	C	$\overline{C}$			In case of contamination, we expect with a high probability (80%) of a major DNA contribution from X, otherwise, the contribution will be minor (20%).
$\Pr(x\_major\mid .)$	0.8	0			
$\Pr(x\_minor\mid .)$	0.2	0			
$\Pr(no\ profile\mid .)$	0				
	Th	$\overline{Th}$			If a third person contribution is involved, we expect to obtain most of the time (80%) a minor DNA profile and otherwise a major DNA profile (20%).
$\Pr(z\_major\mid .)$	0.2	0			
$\Pr(z\_minor\mid .)$	0.8	0			
$\Pr(no\ profile\mid .)$	0	1			
	O		$\overline{O}$		If the watch is not relevant, then we expect to recover no DNA profile from the offender, otherwise, we expect to obtain a major DNA profile (90%) and minor DNA profile (10%) from the offender (which depending on O may be X or an unknown person).
	R	$\overline{R}$	R	$\overline{R}$	
$\Pr(o\_major\mid .)$	0	0	0.9	0	
$\Pr(o\_minor\mid .)$	0	0	0.1	0	
$\Pr(x\_major\mid .)$	0.9	0	0	0	
$\Pr(x\_minor\mid .)$	0.1	0	0	0	
$\Pr(no\ profile\mid .)$	0	1	0	1	
Node 11: binary table					This table is constructed applying four addition rules that dictate the outcome when two contributions from X are mixed together. The following permutative rules have been used: (1) *x_major* + anything else gives *x_major* (2) *x_minor* + *no profile* gives *x_minor* (3) *x_minor* + *x_minor* gives *x_major* (4) *no profile* + *no profile* gives *no profile*
Node 12: binary table					This table determines the outcome in function of the contribution obtained from Node 9, 10, and 11. The four addition rules above are also applied when necessary. For example: If Node 9 = *no profile*, Node 10 = *o_minor* and Node 11 = *x_major*, then the outcome will be *Match_Major/Mix1_Minor.* If Node 9 = *no profile*, Node 10 = *x_minor* and Node 11 = *x_minor*, then the outcome will be *a match,* classified as *any other outcome*. Etc.

Clearly, there are many probabilities that contribute to the overall assessment of the evidence in relation to the Level 3 propositions. Some of these might be within the expertise of the scientist, others relate more to the circumstances surrounding the incident and arrest of the suspect. We see the role of the scientist in such a case as one, not of giving a firm *LR* to the court, more in the role of explaining the sensitivity of the final assessment to the many relevant factors. For the purpose of illustration, we can explore the uncertainties in the *LR* that follow from uncertainties in the assignment of r, c, t, q, and s. As before, these can be treated as unknown parameters and simulations run using Matlab. We assign 80% probability intervals as follows:

r—is the probability that the watch was worn by the offender. Clearly, this is of critical importance to the Level 3 propositions. The watch was found in an area that was locked against public access and the owner of the premises was satisfied that it was not there the night before the incident. On this basis, we assign an 80% interval for r of 0.95 to 0.999.

c—is the probability that DNA from the suspect might enter the process because of contamination during the investigation or at the laboratory. This is a probability that would appear to be within the domain of the scientist. For illustration, we will use an interval of 0.00001 to 0.0005. Ideally, this would be informed by experiments that had been carried out.

t, q—We have seen that the outcome showed the presence of a third unknown profile on that watch. The defence view would be that this profile is that of the offender–who is not

X–whereas the prosecution view will be that DNA from a third person has entered the process. There are some complex issues here that we do not go into. Suffice it to say that, if q were zero then the support for the prosecution proposition would also be zero; on the other hand, if q were one then the prosecution proposition would be well supported by the observed result. These are issues to be explored with the court but for the present sensitivity analysis we will assign a probability interval of 0.2 to 0.8 to t and 0.05 to 0.2 to q. We do not claim that these are in any sense the "right" ranges, but they serve for illustration.

s—The suspect admitted having a watch but there is doubt about whether or nor it was that watch or another one. Evidence has been provided regarding paint splashes visible on the recovered watch and on a photograph showing a watch worn by the suspect. Based on our assessment of this piece of evidence, we expect s to be high, but again this assessment lies more in the domain of the court. For illustration, we will use an interval of 0.85 to 0.99.

In the previous example, we showed the Beta distributions that we used for simulation. Here, we omit the graphs but, for those who might be interested in the detail, we show in Table 8 the Beta parameters for modelling r, c, t, q, and s.

TABLE 8—*80% probability ranges and parameters for the underlying Beta distributions for r, c, t, q and s.*

	Lower Bound	Upper Bound	Alpha	Beta
1-r	0.001	0.05	0.71	34.7
c	0.00001	0.0005	0.055	189.7
t	0.2	0.8	2.06	2.06
q	0.05	0.2	3.39	25.12
1-s	0.01	0.15	1.17	15.84

NOTE: Our fitting algorithm for obtaining the Beta distribution parameters converges efficiently for proportions under 0.5. For that reason the parameters for (1-r) and (1-s) instead of r and s were calculated.

For each simulation, the values r, c, t, q, and s were generated according to their respective distributions and the *LR* was obtained by computing the ratio between the probability of the outcome of interest given O with its corresponding probability given $\overline{O}$. The distribution of the likelihood ratio for the outcome of interest is shown in Fig. 8 and summarised in Table 9.

We repeat that we would see this distribution—and others like it—as a medium for informing discussions with counsel and, if necessary, at court, about the various issues that are relevant to addressing the Level 3 propositions in the case. We emphasise again that the *LR* that we have taken for the sub-level 1 propositions is effectively infinite (though this is not a practice we would support at court). The sensitivity analysis shows a modal value of around 6 for the propositions that would be put to the jury.

Discussion and Conclusion

All those familiar with the presentation of scientific evidence at court will be aware that, particularly in complex cases, the issues put to the jury can descend to a welter of if's, buts, and maybes. We are not claiming that we have an instant solution for resolving such confusion. What we are convinced of, however, is the power of the main themes of our paper as aids to discussing the complicated issues that can interact in cases where small quantities of DNA are brought to the notice of the court. In summary, we emphasis four features of our discussion:

1. The notion of the hierarchy of propositions is a powerful aid for discussions among scientists; between scientists and advocates; and must ultimately play a role in the presentation of evidence

TABLE 9—*Summary statistics for the distribution of the likelihood ratio for the outcome "Match_major/Mix1_minor."*

Minimum	1st Quartile	Median	Mean	3rd Quartile	Maximum
0.298	3.361	5.148	6.587	7.913	96.56

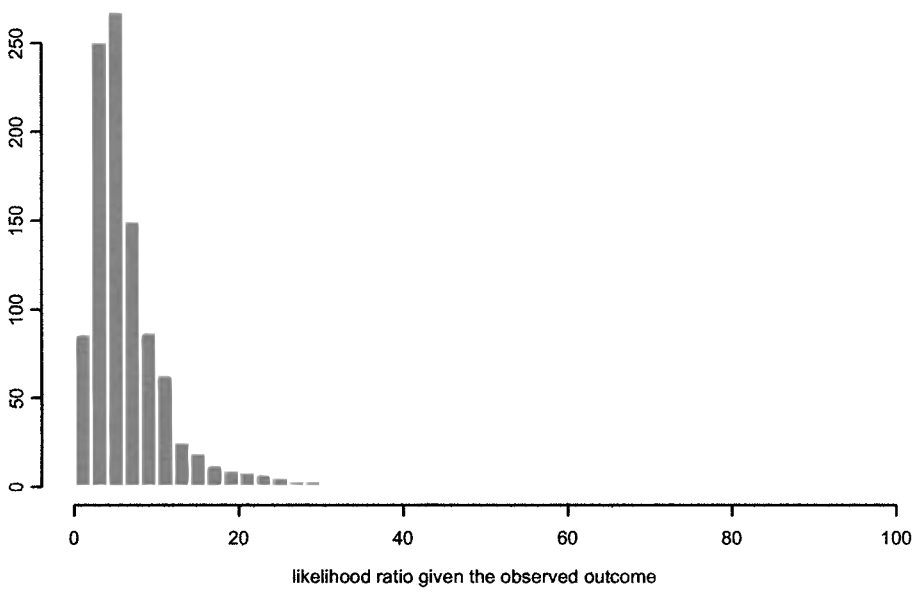

FIG. 8—*Distribution of the likelihood ratio in Case Study 2 for the observed outcomes given the specified uncertainties in r, c, t, q, and s.*

to a jury. The idea that, in relation to any issue, it is essential to address *two* competing propositions is the key to the balanced scientific view.

2. The idea of pre-assessment of a case also contributes to a balanced view. It directs the scientist to consider his/her expectations *before* collecting scientific evidence. To say, after making a particular observation "this is what I would have expected to find" will always invite the suggestion that this is a *post hoc* rationalization.
3. We anticipate that Bayesian networks will play an increasingly important role in forensic science. This does not necessarily mean the vision of the scientist presenting networks at court—but their power in enabling the scientist to understand the fundamental issues in a case and to discuss them with colleagues and advocates is something that has not previously been seen within forensic science.
4. As we said at the outset, the debates relating to the statistics of DNA profiling may have tended to divert attention away from broader issues that are every bit as important. In an individual case, even if the inference with regard to the source of a sample of DNA is effectively indisputable, the inference with regard to whether or not the defendant is the offender may be subject to considerable uncertainty.

We believe that our discussion has important implications for the future of forensic science. A number of the probabilities that enter our analyses are clearly the province of the jury and there may be a view that the scientist should confine him/herself more to the technical level and not stray beyond the source level propositions. This leaves all other interpretative issues to advocacy. However, we believe that the scientist can provide a logical perspective that might otherwise be lacking. We do not attempt to minimize the difficulties but we do believe that the scientist has a role in helping to resolve them.

References

1. Dawid AP, Evett IW. Using a graphical method to assist the evaluation of complicated patterns of evidence. J Forensic Sci 1997;42:226–31.
2. Cook R, Evett IW, Jackson G, Jones PJ, Lambert JA. A hierarchy of propositions: deciding which level to address in casework. Sci and Justice 1998;38:231–40.
3. Cook R, Evett IW, Jackson G, Jones PP, Lambert JA. A model for case assessment and interpretation. Sci and Justice 1998;38:151–6.
4. Cook R, Evett IW, Jackson G, Jones PJ, Lambert JA. Case pre-assessment and review in a two-way transfer case. Sci and Justice 1999;39:103–11.
5. Evett IW, Jackson G, Lambert JA. More in the hierarchy of propositions: exploring the distinction between explanations and propositions. Sci and Justice 2000;40:3–10.
6. Stoney DA. Relaxation of the assumption of relevance and an application to one-trace and two-trace problems. J Forensic Sci Soc 1994;34:17–21.
7. Evett IW. Establishing the evidential value of a small quantity of material found at a crime scene. J Forensic Sci Soc 1993;33:83–6.
8. Evett IW, Weir BS. Interpreting DNA evidence—statistical genetics for forensic scientists. Sunderland: Sinauer Associates, Inc., 1998.

Additional information and reprint request:
Dr. Christophe Champod
The Forensic Science Service
Trident Court
R & D–Interpretation Research Group
Birmingham Business Park
Solihull B37 7YN
Tel: 44 (0) 121 329 5375
Fax: 44 (0) 121 622 2051
E-mail: cchampod@compuserve.com