Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Volcanology and Geothermal Research 128 (2003) 273–285
Journal of volcanology and geothermal research
www.elsevier.com/locate/jvolgeores

Evidence-based volcanology: application to eruption crises

W.P. Aspinall [a,*], G. Woo [b], B. Voight [c], P.J. Baxter [d]

[a] *Aspinall and Associates, Woodside Close, Beaconsfield, Bucks HP9 1JQ, UK*
[b] *Risk Management Solutions, 10 Eastcheap, London EC3M 1AJ, UK*
[c] *Department of Geosciences, Penn State University, University Park, PA 16802, USA*
[d] *Department of Public Health and Primary Care, Cambridge University, Fenner's, Gresham Road, Cambridge CB1 2ES, UK*

Abstract

The way in which strands of uncertain volcanological evidence can be used for decision-making, and the weight that should be given them, is a problem requiring formulation in terms of the logical principles of Evidence Science. The basic ideas are outlined using the explosion at Galeras volcano in Colombia in January 1993 as an example. Our retrospective analysis suggests that if a robust precautionary appraisal had been made of the circumstances in which distinctive tornillo signals were detected at Galeras, those events might have been construed as stronger precursory evidence for imminent explosive activity than were the indications for quiescence, given by the absence of other warning traits. However, whilst visits to the crater might have been recognised as involving elevated risk if this form of analysis had been applied to the situation in January 1993, a traditional scientific consideration of the available information was likely to have provided a neutral assessment of short-term risk levels. We use these inferences not to criticise interpretations or decisions made at the time, but to illustrate how a structured, evidence-based analysis procedure might have provided a different perspective to that derived from the conventional scientific standpoint. We advocate a formalism that may aid such decision-making in future: graphical Bayesian Belief Networks are introduced as a tool for performing the necessary numerical procedures. With this approach, Evidence Science concepts can be incorporated rationally, efficiently and reliably into decision support during volcanic crises.

Keywords: Bayes' Rule; Bayesian Belief Network; decision support; Evidence Science; expert judgment; Galeras volcano; risk assessment; volcanic eruption

1. Introduction

In one of his last volcanology publications, Bruno Martinelli (1997) pointed out that methods to detect and to quantify the dynamics of the processes suspected of triggering eruptive activity are practically non-existent. In this context, and as a physicist, he argued that research should be focused on the physical aspects of volcanic activity, and that more observations and experiments should be performed. It was his view that today's geophysical approach to the investigation of volcanic activity allows, at best, only a static reconstruction of the internal structure of the volcanic complex, and that geological and petrological investigations, as well as the analysis of volcanic gases, while representing substantial sources of information towards accomplishing the task of short-term prediction, are not substitutes for a

* Corresponding author. Tel.: +44-(0)1494-672746; Fax: +44-01494-671239.
E-mail address: willy@aspinall.demon.co.uk (W.P. Aspinall).

0377-0273/03/$ – see front matter
doi:10.1016/S0377-0273(03)00260-9

full examination and understanding of the physics involved.

This is, undoubtedly, a valid viewpoint; volcanology, however, is a science which differs from most others in one crucial respect: crises may arise from time to time which require volcanologists to make immediate recommendations that affect public safety. Their training generally follows the scientific tradition of making observations of Nature, conducting experiments, and developing explanatory models. Where significant uncertainty exists, further research is undertaken to reduce it, in the spirit espoused by Martinelli. But the timescale for such research in volcanology can be very protracted, being dependent, like other observational sciences, on the occasional occurrence of future events (not to mention funding).

The present paper seeks to provide a complementary perspective to Martinelli's contributions to volcanology by describing a formalised risk analysis approach in which such scientific data as are available during a volcanic crisis, albeit inevitably incomplete, insufficient and uncertain, can be used most logically and effectively for assessing and stating hazard levels.

In such circumstances, the application of certain elementary principles for the rational treatment of uncertainty can help in decision-making, especially where responsibility for advice may rest mainly on the subjective judgment of a few scientists or even an individual volcanologist. The basic principles are common to decision-making in many key walks of life, including the law and medicine: for instance, 'Evidence-based Medicine' is a recent idea for the formalised integration of best research evidence with clinical expertise and patient values in medical practice (e.g. Sackett et al., 2000). Similar principles apply to the use of forensic science in the courtroom (e.g. Robertson and Vignaux, 1995) although, in this latter instance, deficiencies in judicial appreciation at appellate level of the rationality of the Bayesian treatment of evidence is currently hindering its full and appropriate use (in Britain, at least – Balding, 1996). The principles are logical rather than mathematical and, as such, should be accessible to the whole volcanological community.

The concepts involved are to be found in an emerging discipline that has been given the overarching title of 'Evidence Science' (Aitken, 1995). Within this paradigm, all aspects of evidential interpretation are treated in a formalised fashion – not just frequentist results from conventional statistical analysis (where these are available), but also uncertain evidence and the weight to be accorded to it, as may be expressed in terms of degree of belief or so-called subjective probability. The fundamental principles are those of probability theory, acting through the precepts of Bayesian statistics (Jeffreys, 1961). The development of a derivative theme in the form of 'Evidence-based Volcanology' (EBV), can be anticipated. For volcanologists involved in decision support for volcanic crisis management and hazard mitigation, genuine practical benefits will be created as this discipline advances. The tragic episode that occurred at Galeras volcano in Colombia in 1993 (Williams, 2001; Bruce, 2001; Monastersky, 2001; Rose, 2001) can be used to illustrate some of the principles involved, and to exemplify the application of these ideas in critical conditions.

2. Galeras, 14th January 1993 – an Evidence Science perspective

Like other Earth scientists, volcanologists tend to interpret evidence in an informal manner, often relying on personal experience and sometimes on 'gut-feeling'. But, with a modest effort, a more structured and logical approach to the use of evidence is possible, which can help arrive at a rational decision.

Suppose that the 1993 Galeras International Decade Volcano Workshop had included a review of some basic principles of Evidence Science, such as Bayes' Rule. This rule, which is a logical theorem, explains how we can invert uncertain information to make general statements on a hypothesis on the basis of special cases. In notational form, it is written as:

$$\Pr(H|E) = \Pr(H)\cdot\frac{\Pr(E|H)}{\Pr(E)}$$

where H is the hypothesis or proposition, E the evidence about H, and the | symbol denotes a conditional probability.

In its more convenient 'odds form' (see e.g. Lindley, 1971), Bayes' Rule shows us how to up-date a hypothesis when new evidence is received:

$$Odds(\mathrm{H}|\mathrm{E}) = \frac{\Pr(\mathrm{E}|\mathrm{H})}{\Pr(\mathrm{E}|\overline{\mathrm{H}})} \cdot Odds(\mathrm{H})$$

where the *odds* is the ratio of the probability of the hypothesis being true to the probability of its being false under the same conditions, the over-line symbol denotes negation, and the ratio of the probabilities of the evidence in the two circumstances, true and false, is referred to as its *likelihood ratio* (LR).

For the present discussion, it is suggested that the threat of an imminent explosion at Galeras could have been considered as a proposition and that, following the acquisition of certain observational data, up-dated or posterior odds of this happening were needed. The posterior odds in this case would be equal to the prior odds of imminent explosion (i.e. before the data were acquired), multiplied by the LR of obtaining the specific evidence, given an explosion is imminent. In equation form:

$$Posterior\ Odds\ (\text{Imminent Explosion}) = \frac{\Pr(\text{Imminent Explosion}|\text{Data})}{\Pr(\text{No Imminent Explosion}|\text{Data})} = \frac{\Pr(\text{Imminent Explosion})}{\Pr(\text{No Imminent Explosion})} \times Likelihood\ Ratio\ (Data)$$

where: *Likelihood Ratio* (*Data*) = Pr(Data | Imminent Explosion)/Pr(Data | No Imminent Explosion).

Jeffreys (1961) graded values of the Bayes' LR in a simple manner, according to multiplicative steps of the square root of 10, with a simple verbal description for the weight of the evidence. A version of his scheme, adapted to the present discussion, is given in Table 1.

A LR close to unity (i.e. approaching a 50:50 bet) indicates very weak or uninformative evidence; high or low values for the LR indicate firm evidence for or against the proposition, respectively.

For the present purposes, a time window of one week is used to define 'imminent' for the Galeras case study. Galeras volcano in Colombia emerged from a 50-yr period of quiescence in 1988 (Cortés and Raigosa, 1997). Based on its subsequent activity over the next few years, the 'base rate' of explosive activity at Galeras might have been estimated as approximately equivalent to one event per year. Thus, the prior odds of an imminent explosive eruption in any one week, as of 13th January 1993, could be taken as:

$$Odds\ (\text{Imminent Explosion}) = \frac{\Pr(\text{Imminent Explosion})}{\Pr(\text{No Imminent Explosion})} = \frac{1/52}{51/52} = \frac{1}{51}$$

Now consider the situation where new data are obtained in the form of distinctive volcano–seismic signals, known as 'tornillos' (Narváez et al., 1997; Gómez and Torres, 1997; Gil Cruz and

Table 1
Grade descriptors of explosion evidence LR values (adapted from Jeffreys, 1961)

$LR > 10^2$	Evidence for imminent explosion is decisive
$10^{3/2} < LR < 10^2$	Evidence for imminent explosion is very strong
$10 < LR < 10^{3/2}$	Evidence for imminent explosion is strong
$10^{1/2} < LR < 10$	Evidence for imminent explosion is substantial
$1 < LR < 10^{1/2}$	Evidence for imminent explosion is just worth a mention
$10^{-1/2} < LR < 1$	Evidence against imminent explosion is just worth a mention
$10^{-1} < LR < 10^{-1/2}$	Evidence against imminent explosion is substantial
$10^{-3/2} < LR < 10^{-1}$	Evidence against imminent explosion is strong
$10^{-2} < LR < 10^{-3/2}$	Evidence against imminent explosion is very strong
$LR < 10^{-2}$	Evidence against imminent explosion is decisive

Chouet, 1997). Tornillos are a sub-class of so-called 'long period' volcano–seismic signals that are postulated to be indicators of unsteady fluid transport where there is high impedance between fluid and solid (Chouet, 1996). In some instances, particularly as shown by studies at Galeras since January 1993, their occurrence has been interpreted as a predictive marker of impending eruptive activity (Gil Cruz and Chouet, 1997). The LR for this particular phenomenon, given an explosion is imminent, is:

$$LR(\text{Tornillos}) = \frac{\Pr(\text{Tornillos}\,|\,\text{Imminent Explosion})}{\Pr(\text{Tornillos}\,|\,\text{No Imminent Explosion})}$$

For a sceptical scientist, a tornillo might be downplayed as an abstruse seismological signal of questionable origin and little predictive significance, i.e. its LR might be very near unity. However, for a convinced seismologist, the tornillo might have been perceived more importantly – as a signal of potentially high predictive power. But just how high? Tornillos are a relatively unusual feature of volcano–seismic activity anywhere, having been recorded only at a few active volcanoes (see e.g. Gómez and Torres, 1997).

Dedicated seismic network monitoring of Galeras commenced in February 1989, when fumarolic activity was increasing (Cortés and Raigosa, 1997). Seismological recordings obtained in July 1992, after the emplacement of an andesitic lave dome in the main crater in 1991, provided the first evidence that the build-up to a dome-destroying vulcanian explosion at Galeras might sometimes involve a tornillo-generating precursor: from 11th July, nine tornillo signals were observed, with four of these occurring just before the significant explosion of 16th July 1992. While a sceptic might have viewed this single episode and its timing as merely accidental, the fact is there was a paucity of experience with which to assess a conjectured link between tornillos and dome-destroying explosions at Galeras. Thus, in January 1993, all that was known was that tornillos had preceded the one dome-destroying explosion which had been adequately monitored (16th July 1992).

In these circumstances, a very cautious attitude to adopt would have been to assume that it was very likely that any immediate subsequent explosion of the Galeras dome would be preceded by this unusual phenomenon. For most scientists, however, taking a firm position based on a single instance runs contrary to inclination and training, and many would have had difficulty ascribing a strong likelihood to the proposition that there might be a reliable association between tornillos and explosions. A risk analyst, on the other hand, might have felt obliged to place a high probability in the numerator of the LR, under the conditions then prevailing at this particular volcano, until evidence was forthcoming that contradicted the precautionary hypothesis. For argument here, let us assume he would have ascribed a high probability that the association could hold true in the short term, say Pr(Tornillos | Imminent Explosion) = 0.9 (i.e. approximately *10:1 on*, in betting odds terms, a relatively high diagnostic test sensitivity). The scientist in contrast, if pressed, might have put the chance of this association as high as 'evens' in his terms (i.e. 'I don't know one way or the other') but, given the lack of a recognised precedent from other eruptions, is more likely to have suggested much lower odds in favour of the proposition: let us say he put theses odds at 1/4, in other words: Pr(Tornillos | Imminent Explosion) = 0.2.

For the denominator of the LR, the history of all local tornillo observations, albeit short, allows an estimate of the tornillo 'false alarm' rate (or diagnostic test specificity) to be made. In the present case, this is the chance of having a week or more in which one or more tornillos occur daily but are not followed promptly by an explosion. At Galeras, a few tornillos were recorded in March 1989, and a singleton in August 1992, but no eruptive activity followed on either occasion. Thus, as far as was known in early 1993, tornillo episodes at the volcano were comparatively rare, and a false alarm rate of perhaps once to three times a year might have been judged reasonable, in the light of the experience up to that time. Let us take Pr(Tornillos | No Imminent Explosion) = 2/51 = 0.0392, as our value for the denominator of the LR, *LR*(Tornillos), (there being 51 weeks each year without an explosion in two of which tornillos occur, on average). This would make the value of the overall LR for the scientist equal to $0.2 \div 0.0392 = 5.1$, using the chosen probabilities in favour of occurrence association and the weekly false alarm rate, respectively. For our imaginary risk analyst, the value of this LR would be 22.96 (i.e. $0.9 \div 0.0392$).

Thus, on Jeffreys' verbal scale, while the weight of the tornillo evidence for an imminent explosion of the Galeras dome in January 1993 (given the unique association with the 16th July 1992 explosion, and no other monitoring information) would have been *substantial* for the scientist, it would have been *strong* for the risk analyst. So, based solely on the observation of tornillos occurring in the week or so prior to the field trip, the hypothetical scientist's posterior odds for an imminent explosive eruption might have been brought up from the background base-rate of 1/51 to something close to 1/10 ($1/51 \times 5.1$), if Bayes' Rule had been applied to this line of evidence; for the risk analyst, the posterior odds would have been updated to about 1/2.2 ($1/51 \times 22.96$), equivalent to Pr(Imminent Explosion) = 0.31.

At Galeras, however, there was other monitoring evidence to take into account (Narváez et al., 1997), with the potential to modify the posterior odds of imminent explosion indicated by the tornillo evidence.

3. Adding other strands of evidence

Solely for the purposes of the present discussion, we consider now three additional 'traits' that might have been supposed by the workshop participants to be indicative of a forthcoming explosive eruption. For instance, these other traits which, for simplicity, can be assumed to be generally independent of the tornillo trait, might be denoted by:

Trait S: anomalous Seismic activity (other than tornillos)

Trait G: elevated Gas flux and/or temperature

Trait D: increased Deformation

Monitoring of these physical parameters was being pursued at Galeras prior to 14th January 1993, but it seems none was manifest significantly or anomalously (Narváez et al., 1997).

Although no gas flux measurements had been possible since mid-December, nothing anomalous about gas emissions was noted at the time of the workshop; seismicity was generally low, and no significant deformation was detected. There was awareness amongst local scientists that whilst seismic, gas and deformation levels had reduced since January 1992, an explosion had occurred in July 1992, nonetheless. There was also some appreciation at the workshop that the low seismicity and gas output might reflect sealing of the system (Stix, 1993), and this idea may have received an airing among some delegates. In the circumstances, had any one of the three traits been perceived as abnormal, then concerns about safety in the crater may well have been amplified. Whilst the scientists involved were aware of the volcano's explosive potential – and, on that account, sought to limit the numbers exposed in the crater – the contemporaneous absence of any anomalous behaviour in *all* of these traits might have been taken as a comforting sign by some.

In these circumstances, consider the LR, as of 14th January 1993, for the joint absence of all three traits:

$$LR(\text{No Trait}) = \frac{\Pr(\text{No Trait S, G, D} \mid \text{Imminent Explosion})}{\Pr(\text{No Trait S, G, D} \mid \text{No Imminent Explosion})}$$

To elaborate this numerically, we assume, on the basis of experience of other volcanic eruptions, that enough confidence might have been placed in the precursory value of any one of these traits, if it were present, for the numerator to be estimated between 1/5 and 1/2. To illustrate the meaning of the relevant probabilities, an optimistic proponent might have posited that, at the time in question, 80% of all imminent eruptions of a volcano like Galeras would be preceded by some evidence of anomalous behaviour in one or more of these traits. A less confident person might have felt that, say, only 50% of imminent eruptions are preceded by some discernible evidence of this kind – for this numerical illustration, however, we adopt the 80% diagnostic sensitivity value for these other traits, i.e. Pr(No Trait S, G, D | Imminent Explosion) = 0.2.

The denominator of *LR*(No Trait) is a measure of the absence of false alarms, and can be quantified by evaluating [1−Pr(Trait S, G, D| No Imminent Explosion)]. If we assume that in an average year at Galeras there might be a number of different weeks (perhaps 2–4) when one or other of these traits was sufficiently anomalous to lead to anxiety about imminent eruptive activity, later seen to be unwarranted, the denominator of *LR*(No Trait) could be given a value between 1–4/51 and 1–2/51, say 48/51 (we do not claim such numbers are correct, necessarily, but adopt them here to allow the method of analysis to be demonstrated). The LR for the absence of these three traits, assuming the higher diagnostic sensitivity, would then be numerically equivalent to 51/240 ($1/5 \div 48/51$), or 0.2125. In Jeffreys' verbal terms, the weight of the evidence that there was an absence of other precursory traits, which would militate against an imminent explosion, could be ruled *substantial* (but only just, being numerically fractionally more than *worth a mention*).

In the case of Galeras, therefore, the additional evidential likelihood factor, obtained by recognising that even though the unusual tornillo signals were present there was also an absence of other precursory traits, would have served to reduce the gain in the scientist's posterior odds for an imminent explosive eruption from about 1/10 (judged on his interpretation of the evidence of the tornillos alone) to about 1/47 ($1/10 \times 51/240$); put another way, the probability of an explosion at Galeras within a week could have been assessed by such a scientist, on or just before 14th January 1993, at 2.1% if all the main strands of volcanological evidence had been weighed together by Bayes' Rule (cf. the prior base-rate odds estimate of 1/51 per week, or probability $\sim$1.9%). In other words, with this perspective, the one strand of evidence would have balanced out, almost exactly, the other.

If caution were paramount, the analyst might have felt that he should accept that only 50% of imminent eruptions are likely to be preceded by discernible trait changes – i.e. Pr(No Trait S, G, D|Imminent Explosion) = 0.5, and the No Trait LR would then be set to 51/96 ($1/2 \div 48/51$), or 0.53. In Jeffreys' verbal terms, the weight of this strand of evidence, directed against an imminent explosion, would be *worth a mention*, but is not *substantial*. As a consequence, the analyst's joint probability of an imminent explosion, would have been revised down from 31% to 19% (i.e. posterior odds = $1/2.2 \times 0.53$) when both sets of evidence were taken into account.

Thus, this retrospective analysis suggests that in a precautionary approach, the affirmative power of the tornillo evidence for an imminent explosion could have been given greater weight than the negative evidence provided by the absence of other, more conventional precursors. Without a formalised quantitative procedure of this kind for combining the available evidence, finding a balance between unusual observations, such as the tornillos on the one hand, and alternative sources of evidence on the other, can sometimes become highly contentious amongst volcanologists of different scientific backgrounds. This difficulty is not new in volcanology, and affects both societal and personal risk assessment.

4. Individual decisions on risk exposure

At a personal level, and with appropriate data available, an individual volcanologist should be able to make a balanced posterior estimate of the odds of his or her getting injured, in a manner like that just described.

However, in a situation such as that at Galeras in January 1993, three different personal responses can be envisaged:

(1) those sceptical of the information gain of *any* precursory traits might have felt inclined to stick with the prior odds against an explosion in any one week of 1/51;

(2) those sceptical of the significance of tornillos and their interpretation, but reassured by the absence of the more traditional indicators, might have calculated relatively reassuring posterior odds for an imminent explosion of only 1/240 ($1/51 \times 51/240$);

(3) those suspicious of the validity of the tradi-

tional indicators, but putting strong reliance on the diagnostic implications of tornillos, could have arrived at the much more alarming posterior odds in favour of an imminent explosion of 1/2.2, i.e. more than $100\times$ higher than (2).

For a single 4-h trip into the crater (i.e. 1/42 of a week), these latter weekly odds equate to a short-term chance of a life-threatening event occurring during the visit of about 1/92. Any individual volcanologist spending 24 h in circumstances where this level of threat is present would be taking the same magnitude of risk as an astronaut in a space launch (Wilson and Crouch, 2001, table 7-1A); and one who repeated 16 or more such excursions would have a less than 50% chance of surviving through his career unscathed.

Offsetting the hazard exposure implied by the tornillos in the Galeras case by the evidence of the absence of other precursory traits reduces the risk odds for one 4-h visit to 1/174; while this is a more acceptable risk exposure, it still entails a less than 50% chance of survival for anyone who accumulates 500 h of time spent on such visits. In either case, the inferred risks of serious injury appear significantly higher than would be voluntarily accepted in almost any other professional activity.

On the other hand, those espousing position (2), assigning no significance to the tornillos, would have been embarking on the expedition with perceived short-term odds against an untoward outcome of about 1/10,000 ($1/240 \div 2$) – with these odds, an individual would have the expectation of a 99% chance of surviving 100 randomly-timed visits.

Although it is not known what the typical risk tolerance spectrum is of volcanologists who work on active volcanoes (and this is something that would be desirable to ascertain), for the scientists who went on the ill-fated field trip to the crater of Galeras on 14th January 1993 (none of whom, apparently, would have been strongly persuaded by the tornillo evidence as it was perceived at the time), the chances of an imminent explosion would have appeared to be very low. Thus, with the benefit of hindsight, it would be invidious to criticise an expedition leader for taking with him or her other scientists who were keen to enter the crater and whose posterior odds of imminent explosive activity were probably 'professionally acceptable'.

While such volcanologists may decide for themselves, as individuals, whether the risk posed by a given situation on a volcano is tolerable or not, public safety requires a more conservative view of risk for other people. In the Galeras example, had these Bayes' Rule calculations been done at the time, the upper bound on possible posterior odds of an imminent explosion might have been judged a sufficient basis for advising the authorities to issue a public warning and to bar non-essential personnel from approaching the crater. When, subsequently, the significance of tornillos at Galeras became better established in a series of further explosions, local authorities were apprised in May 1993 of a likely impending eruption, and the anticipated explosion duly occurred on 7th June 1993 (Narváez et al., 1997). This successful warning was achieved because, in the six explosions in 1992–1993, all but one were preceded by tornillos, giving an overall association rate of 5 out 6, and the independence of the sixth from tornillos is debatable; in other words, post hoc odds of 5/6 (i.e. *5:1 on* in betting parlance) would be obtained for a positive association during that particular phase of the volcano's activity. In the light of this, the probability of 0.9 for association, that was assumed above in our analysis of the LR for the hypothetical risk analyst, can be seen to be not grossly conservative for the conditions at Galeras at the time.

The question of what level of exposure should be accepted by volcanologists, and the impact that misjudgments of personal risk have on professional credibility, are topics for another forum: a related challenge, however, is to find a practical approach that addresses the problems of the interpretation of evidence for volcanic hazard and risk estimation when that evidence has associated with it significant scientific uncertainty.

5. Formulating a practical procedure

No matter how much scientific data are gathered by volcanologists, or how well and promptly they are analysed, uncertainty over eruption forecasting will always remain (Woo, 1999). Hence there is the need for a structured and logical approach to interpreting the evidence available and to decision-making in volcanic crises, such as when to advise the raising of a red alert, for instance.

For decision or prediction issues of this kind, that involve reasoning with uncertainty, the generalised Bayesian Belief Network (BBN) is an increasingly accepted formalism for determining a rational choice or for providing decision outcome solutions. This routine provides an automated means for performing the calculations involved in the application of Bayes' Rule to cases such as that just discussed in respect of Galeras. The BBN itself is a graphical construct in which multiple uncertain variables are represented by separate nodes, and causal or influence links between nodes are represented by arcs (Jensen, 1996). Associated with each node is a set of conditional probability values, expressing the relationships of the states of that node to any others in the network to which it is linked. These relationships can be given in terms of statistical probability distributions, when data are plentiful, as discrete condition states when hard information is available, or as subjective probabilities or expert opinion when evidence is uncertain and sparse. Behind the graphical interface of a BBN lies the numerical means for computing possible Bayes' Rule outcomes, with whatever type of information is input.

Thus, once a complete BBN has been constructed, it can be executed quickly and efficiently using an appropriate propagation algorithm to calculate the full joint probability table for all factors in the whole model. There are two key features that make BBNs attractive for EBV: (1) new observations can be entered on the relevant node, and their effect propagated through the net in any direction, immediately updating the marginal distributions of all nodes to which

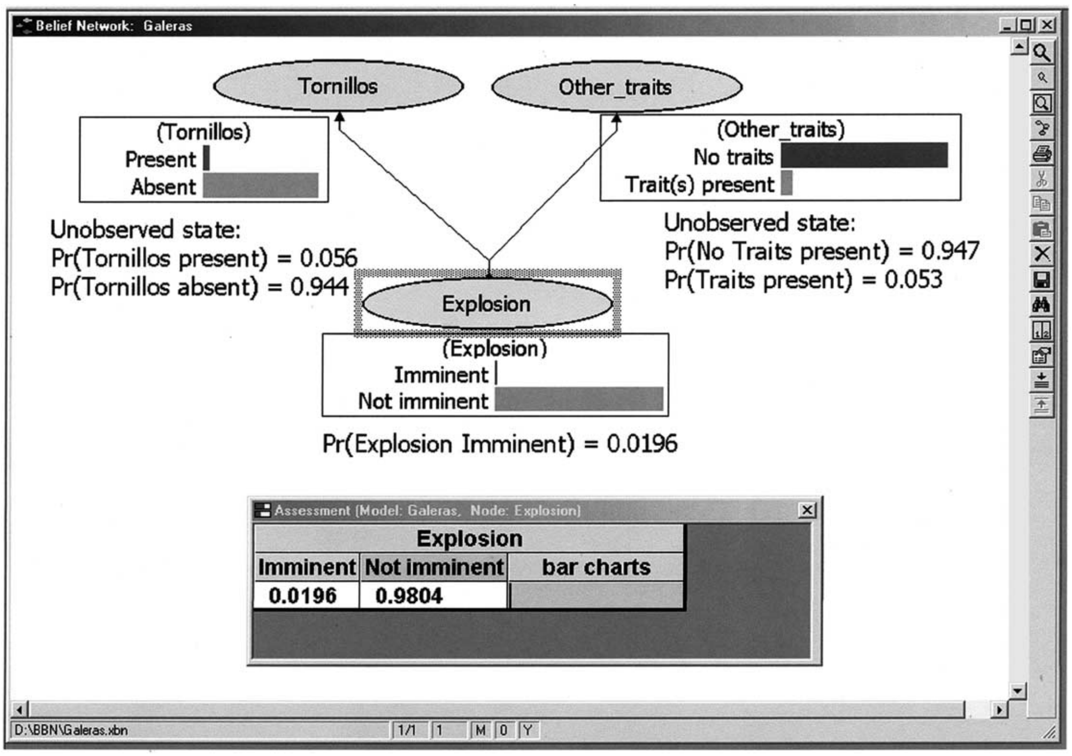

Fig. 1. BBN showing prior probability of imminent explosion (Pr = 0.019), when there is no evidence available about *Tornillos* or *Other_Traits*.

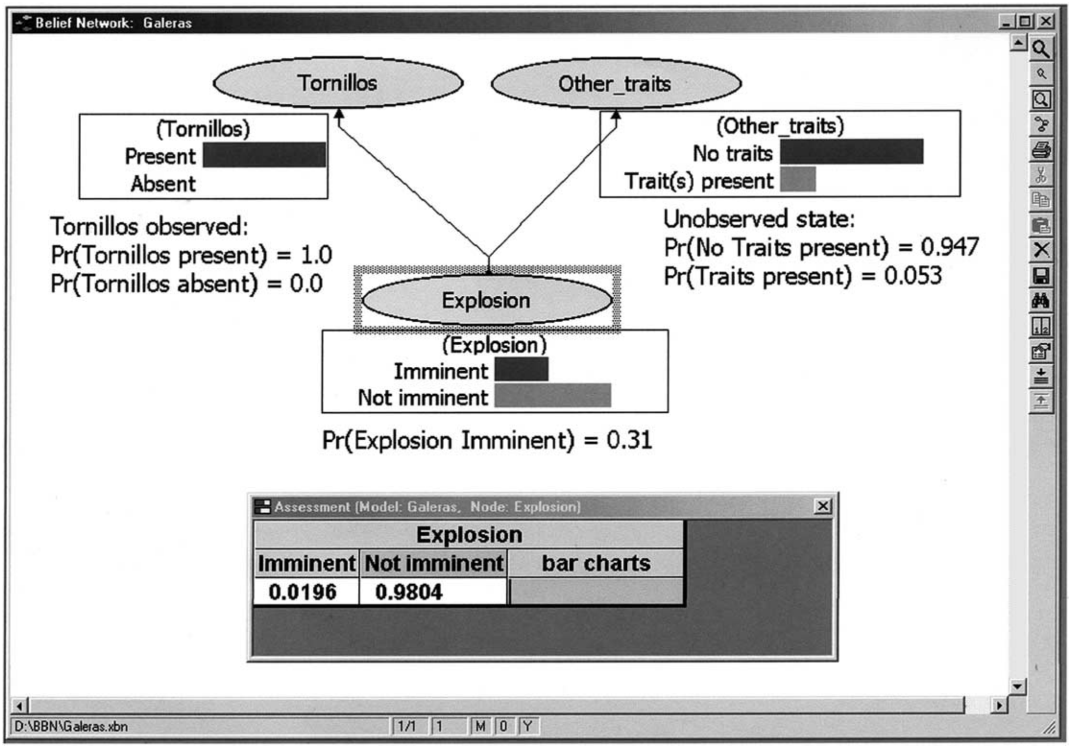

Fig. 2. BBN showing updated probability of Explosion (Pr = 0.32) if *Tornillos* are recorded, but evidence about *Other_Traits* is not available.

a link exists; and (2) by implementation of Bayes' Rule, a BBN model can be used for inference as well as forward prediction.

Here, the way in which a BBN might be applied in an evidence-based science approach to the Galeras tornillo issue is briefly described. Fig. 1 shows an elementary network, depicting the two factors (Tornillos, and 'Other_traits') that could be used as evidence for inferring the probability of an imminent 'Explosion'. The inset window is an interactive feature that allows the user to enter the prior odds for an imminent explosion; similar windows can be opened to enter conditional odds for recording tornillos and for their false alarm rate, prompting numerical calculation of the relevant LR by Bayes' Rule, or for any other factor (e.g. 'Other_traits'). In this implementation, the boxes under each node show the prior probabilities for each factor as bar charts and as numerical values, BEFORE observations become available, so the computed probability Pr(Imminent Explosion) = 0.0192, which is just the prior base rate (odds 1/51), as discussed above.

If, now, tornillos are actually recorded, the condition of the tornillos node switches from 'Unobserved' to 'Present', and the inference as to the probability of an imminent Explosion changes, as shown in Fig. 2. The revised value obtained, Pr(Explosion Imminent) = 0.31, asserts the posterior probability of an explosion (within the following week), given this definite new evidence (and other assumed odds for false alarm rate, etc.). At this stage, the status of evidence about 'Other_traits' is yet to be entered.

In the situation at Galeras in January 1993, the other prognostic traits that might have signaled imminent explosive activity (e.g. anomalous trends in seismicity, gas flux, deformation) were reportedly absent. This additional evidence, concerning the likely condition of the volcano, is incorporated in the BBN in Fig. 3 by setting the node state for 'Other_traits' to Absent: as a re-

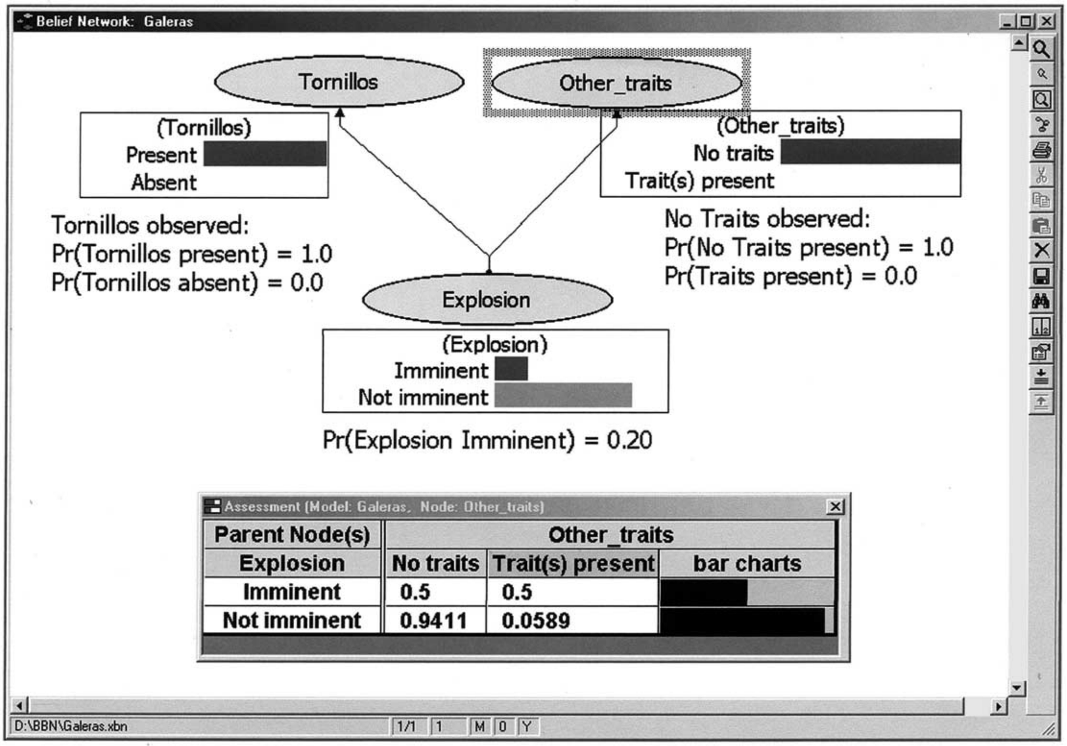

Fig. 3. BBN showing revised updated Explosion probability (Pr = 0.2) given now that *Tornillos* are recorded *and Other_Traits* are NOT present

sult, the inferred probability of explosion, Pr-(Explosion Imminent) = 0.2, remains still elevated, because of the explicit strength of the tornillo evidence interpretation, but much less so than if that evidence had been used in isolation.

In this example, just two separate strands of evidence are involved, connecting to one outcome node ('Explosion'), so the degree to which each contributes to the inferred probability of that outcome might be easily estimated, or even guessed at, from the prior odds assigned to each factor. The BBN formulation, however, provides an easy and convenient way for the exact probability values to be calculated reliably.

In much more complex situations, such as would obtain in a major eruption, the problem of first assembling, then tracking all the compound interactions of multi-dimensional multivariate uncertainties is almost intractable without recourse to a tool such as the BBN. It should not be surprising then if hurried pronouncements made on the spur of the moment in such circumstances, without proper consideration of the relative importance of all the elements of evidence, turn out to be fallacious in the light of subsequent events.

Turning to the current eruption of the Soufrière Hills volcano on Montserrat as a case in point, Fig. 4 illustrates typical elements for just such a situation: estimating the risk of a major collapse of the growing dome when there is evidence from recent experience that this may be triggered by torrential rainfall.

Other factors, however, also come into play in a variety of ways, and this BBN model shows how they might be accommodated. For instance, it is more likely that torrential rain will fall during the rainy season than in the dry season, so the time of year ('Season') is a factor influencing the probability of collapse. However, a weather forecast may be available that predicts heavy rain will fall on the mountain today (e.g. for 'Heavy Rainfall' on the Forecast node Pr = 1 on Fig. 4), or

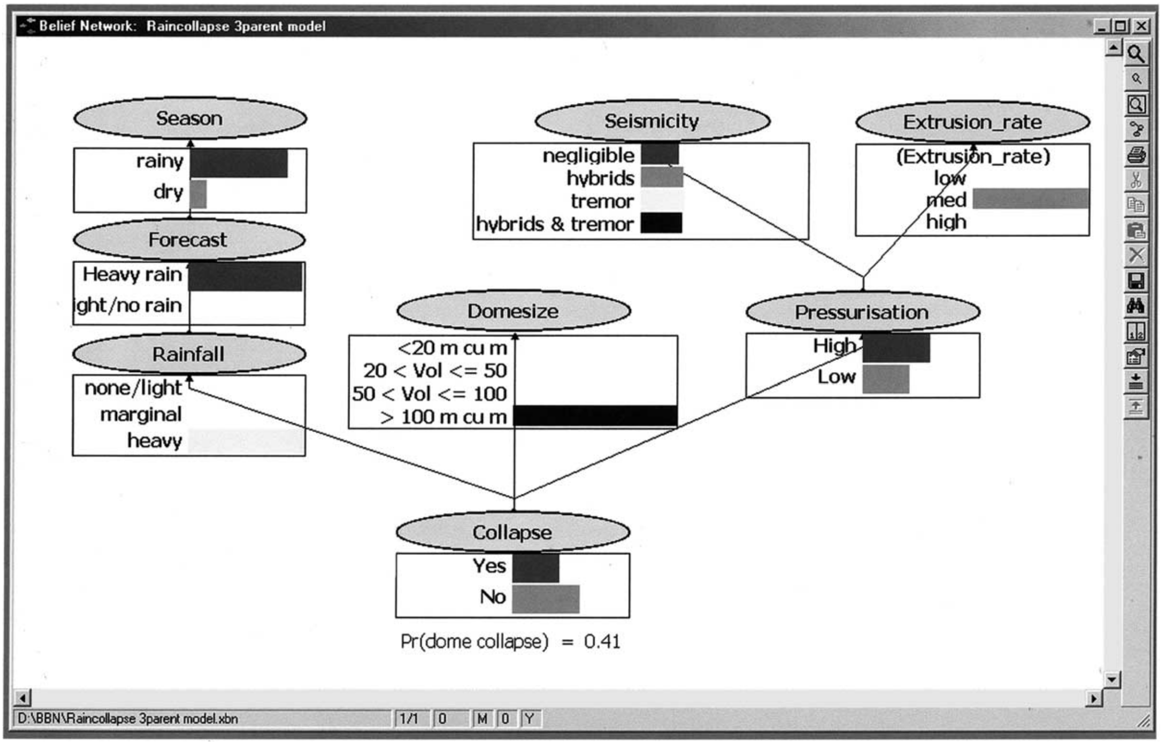

Fig. 4. Example of a BBN for assessing the probability of a dome collapse in a situation with multiple evidential factors, interacting in complex ways (see text for discussion).

that the weather will be fine. In these circumstances, the information about the season is now redundant as far as collapse risk is concerned, and the structuring of the BBN takes care of this when calculating the odds of a collapse – in BBN parlance, the factors 'Season' and then 'Forecast' are successively 'explained away' by the 'Forecast' and 'Rainfall' evidence, respectively. (This analysis could be extended to include a conditional probability for actually getting the defined amount of rainfall, given an evaluation of the reliability of local weather forecasts!)

In Fig. 4, two other factors are depicted that could influence dome collapse risk: the size of the dome, 'Domesize', and the state of its internal 'Pressurisation'. Clearly, there is little chance of a big collapse if the dome is very small, and the risk of a major avalanche increases as the size of the dome grows. This factor can be quantified in terms of volume, and discretised into a few representative categories, as shown in Fig. 4. It is usually straightforward to see directly how big the dome has become, and the corresponding state of the 'Domesize' node, as an observable, is easily determined. Internal pressurisation, on the other hand, is not amenable to direct observation or measurement, and has to be inferred from other indicators, such as tilt, seismicity or extrusion rate. It is therefore a 'latent variable' in this model and its influence on the assessment of probability of collapse can only be inferred, depending on the weight that is given to related observables, such as 'Seismicity' and 'Extrusion_rate' in Fig. 4, if and when these data are available.

For such a situation, it is not immediately obvious how the interplay between all the various evidential factors could influence the outcome probability, and a fully structured model is essential for testing all possibilities. It is much easier, and more reliable, to solve this by aggregating a number of individual factors, each separately assessed in terms of its own LR, than to attempt an overall global estimate by judgment or guesswork.

Thus, the BBN procedure can properly take account of diverse strands of evidence in very complex relational structures, and the formulation is capable of accommodating and harmonising information from variables that are measured or change on different timescales. The individual elements of evidence available may be in one of two basic forms: categorically definite ('hard') information (as in an exceedance over a threshold, for instance), or observed or measured but uncertain ('soft') data. For scientists monitoring volcanic activity, however, the nature and content of the evidence can change as circumstances change, and some lines of data may become intermittent, unobservable or unavailable (e.g. instrument failure; cloud cover), representing major difficulties for interpretation in terms of hazard estimation. The BBN formulation offers a means for responding efficiently to this challenge, which, otherwise, is difficult to describe in any quantitative sense. Furthermore, the BBN can accommodate elements of negative evidence that may have potentially serious implications (e.g. when something stops happening, such as a sudden shutdown of gas output during a dome-building episode). With such a framework in place, the sensitivity of the assessed probability of an event to any number of such varying and variable factors can be resolved objectively in terms of an evidence-based decision.

6. Summing up

'Evidence-based Volcanology' is emerging as a formal, idiomatic model for handling uncertain scientific information in decision support activities for volcanic emergencies. Within this concept, the Bayes' Rule LR is the correct, logical way to present the weights of different pieces of evidence. It has two overwhelming practical advantages: the first is that it is potentially capable of dealing with all situations, however complex, whereas a traditional frequentist approach will only produce a correct answer where single propositions are tested against one piece of evidence, one at a time; the second is that evidence expressed in the LR form can logically be combined with other evidence. The use of frequencies requires an intermediate step in reasoning: significance tests of assertions cannot logically be combined with other evidence at all, particularly where that additional evidence may take the form of expert judgment or other subjective probability.

What we have described above is a formalised probabilistic procedure for combining various kinds of precursory evidence and, in particular, a way for deciding how the influence of unusual information (e.g. the tornillo observations) might be weighed. Clearly, the odds that could be used for the various factors in the Galeras case are endlessly debatable. Without such a procedure, however, deciding what balance should be struck between different forms of evidence has to depend upon what is, in essence, a subjective exercise of pattern and behaviour recognition: sometimes, as with the periodic destruction of a succession of domes at Mount St. Helens, 1981–1986 (Swanson and Holcomb, 1989), and in the 1997 sequence of repetitive eruptive events at the Soufrière Hills volcano, Montserrat (Voight et al., 1998), this traditional approach can be successfully used for forecasting. But, in other circumstances, it may generate irreconcilable and unhelpful contentions amongst scientists of different persuasions, as happened in the notorious controversy concerning the volcanic crisis on Guadeloupe, in 1976 (Fiske, 1984).

In practical terms, the graphical BBN provides a framework for undertaking the routine processing of all the elements of evidence that can, and should, be considered collectively for a proper, defensible use of scientific information and scientific opinion in decision-making during a volcanic crisis. With suitable software, a working system to implement the BBN can be put into place quite rapidly and, once set up, volcanic eruption probabilities can then be updated regularly and easily by observatory staff. These updates would be of practical value in facilitating decisions on daily alert levels (Aspinall and Cooke, 1998), for instance, and judgments as to whether or not to undertake fieldwork.

This structured, evidence-based approach can be coupled with other formalised procedures for the elicitation of expert opinion (e.g. Aspinall and Woo, 1994), and with probability tree representations of plausible hazards (e.g. Newhall and Hoblitt, 2002), to provide a complete auditable trail for the way scientific information has been used. It thus offers an invaluable means for ensuring that rational decisions are made and that sound advice is given when confronting an eruption.

Acknowledgements

The authors are very grateful to Roby Carniel for his meticulous and helpful review and suggestions for improving the manuscript. This work has been undertaken partly with the support of the Montserrat Volcano Observatory (funded through the British Geological Survey by the UK Government Department for International Development), and partly with support from the EC Project MULTIMO (ECG1-2000-00574). W.P.A. also acknowledges a Benjamin Meaker Visiting Fellowship at the Institute of Advanced Studies, Bristol University. The Microsoft Research evaluation tool MSBNx® v1.4.2 was used to implement the BBN trials. The views expressed are those of the authors, and not necessarily of any organisation with which they may be associated.

References

Aitken, C.G.G., 1995. Statistics and the Evaluation of Evidence for Forensic Scientists. Wiley, Chichester, 276 pp.

Aspinall, W., Cooke, R.M., 1998. Expert judgment and the Montserrat Volcano eruption. In: Mosleh, A., Bari, R.A. (Eds.), Proceedings of the 4th International Conference on Probabilistic Safety Assessment and Management PSAM4, September 13–18, 1998, New York. Springer, London, vol. 3, pp. 2113–2118.

Aspinall, W.P., Woo, G., 1994. An impartial decision-making procedure using expert judgment to assess volcanic hazards. In: 'Large Explosive Eruptions', Accademia Nazionale dei Lincei/British Council International Symposium, Rome, 24–25 May 1993. Atti dei convegni Lincei 112, pp. 211–220.

Balding, D., 1996. DNA profiles and the courts. R. Statist. Soc. News 24, 1–2.

Bruce, V., 2001. No Apparent Danger. Harper Collins, New York, 239 pp.

Chouet, B.A., 1996. Long-period volcano seismicity: Its source and use in eruption forecasting. Nature 380, 309–316.

Cortés, G.P., Raigosa, J., 1997. A synthesis of the recent activity of Galeras volcano, Colombia: Seven years of continuous surveillance, 1989–1995. J. Volcanol. Geotherm. Res. 77, 101–114.

Fiske, R.S., 1984. Volcanologists, journalists, and the public: A tale of two crises in the Eastern Caribbean. In: Explosive Volcanism: Inception, Evolution, and Hazards. Natl. Res. Council (US), Studies in Geophysics, National Academy Press, Washington, DC, pp. 170–176.

Gil Cruz, F., Chouet, B.A., 1997. Long-period events, the most characteristic seismicity accompanying the emplacement and extrusion of a lava dome in Galeras Volcano, Colombia, in 1991. J. Volcanol. Geotherm. Res. 77, 121–158.

Gómez, D.M., Torres, R.A., 1997. Unusual low-frequency volcanic seismic events with slowly decaying coda waves observed at Galeras and other volcanoes. J. Volcanol. Geotherm. Res. 77, 173–193.

Jeffreys, H., 1961. Theory of Probability. Oxford University Press, Oxford, 459 pp.

Jensen, F.B., 1996. An Introduction to Bayesian Networks. Springer, New York, 178 pp.

Lindley, D.V., 1971. Making Decisions. Wiley, London, 207 pp.

Martinelli, B., 1997. Volcanic tremor and short-term prediction of eruptions. J. Volcanol. Geotherm. Res. 77, 305–311.

Monastersky, R., 2001. Under the volcano. Chronicle of Higher Education, March 30.

Narváez, L. et al., 1997. 'Tornillo'-type seismic signals at Galeras volcano, Colombia, 1992–1993. J. Volcanol. Geotherm. Res. 77, 159–171.

Newhall, C.G., Hoblitt, R.P., 2002. Constructing event trees for volcanic crises. Bull. Volcanol. 64, 3–20.

Robertson, B., Vignaux, G.A., 1995. Interpreting Evidence: Evaluating Forensic Science in the Courtroom. Wiley, Chichester, 240 pp.

Rose, W.I., 2001. Volcanic irony – hard lessons from the dark side of volcanology (book reviews). Nature 411, 21–22.

Sackett, D.L., Straus, S.E., Richardson, W.S., Rosenberg, W., Haynes, R.B., 2000. Evidence-based Medicine. Churchill Livingstone, London, 261 pp.

Stix, J., 1993. Report of trip to Galeras Volcano, 28 July–16 August 1992. Abs. 1993 Galeras International Decade Volcano Workshop, Pasto, Colombia.

Swanson, D.A., Holcomb, R.T., 1989. Regularities in growth of the Mount St Helens dacite dome, 1980–1986. In: Fink, J.H. (Ed.), IAVCEI Proceedings in Volcanology: Lava Flows and Domes, vol. 2, Springer, Berlin, pp. 1–24.

Voight, B. et al., 1998. Remarkable cyclic ground deformation monitored in real time on Montserrat and its use in eruption forecasting. Geophys. Res. Lett. 25, 3405–3408.

Williams, S., 2001. Surviving Galeras. Little, Brown, London, 270 pp.

Wilson, R., Crouch, E.A.C., 2001. Risk-Benefit Analysis. Harvard University Press, 370 pp.

Woo, G., 1999. The Mathematics of Natural Catastrophes. Imperial College Press, London, 292 pp.