JOURNAL OF GEOPHYSICAL RESEARCH, VOL. 114, B03213, doi:10.1029/2008JB005908, 2009

Principles of volcanic risk metrics: Theory and the case study of Mount Vesuvius and Campi Flegrei, Italy

Warner Marzocchi[1] and Gordon Woo[2]

Received 2 July 2008; revised 16 December 2008; accepted 14 January 2009; published 28 March 2009.

[1] Despite volcanic risk having been defined quantitatively more than 30 years ago, this risk has been managed without being effectively measured. The recent substantial progress in quantifying eruption probability paves the way for a new era of rational science-based volcano risk management, based on what may be termed ''volcanic risk metrics'' (VRM). In this paper, we propose the basic principles of VRM, based on coupling probabilistic volcanic hazard assessment and eruption forecasting with cost-benefit analysis. The VRM strategy has the potential to rationalize decision making across a broad spectrum of volcanological questions. When should the call for evacuation be made? What early preparations should be made for a volcano crisis? Is it worthwhile waiting longer? What areas should be covered by an emergency plan? During unrest, what areas of a large volcanic field or caldera should be evacuated, and when? The VRM strategy has the paramount advantage of providing a set of quantitative and transparent rules that can be established well in advance of a crisis, optimizing and clarifying decision-making procedures. It enables volcanologists to apply all their scientific knowledge and observational information to assist authorities in quantifying the positive and negative risk implications of any decision.

Citation: Marzocchi, W., and G. Woo (2009), Principles of volcanic risk metrics: Theory and the case study of Mount Vesuvius and Campi Flegrei, Italy, *J. Geophys. Res.*, *114*, B03213, doi:10.1029/2008JB005908.

1. Introduction

[2] Volcanic hazard and risk studies have a prominent role, both for society and volcanology itself, being the principal interface between science and public policy. Nevertheless, volcanic risk has been managed without being effectively measured. The main reason lies in the scarce knowledge, extreme complexity, nonlinearity, and large number of barely observable degrees of freedom of a volcanic system. These sources of uncertainty make deterministic prediction of the evolution of volcanic processes difficult if not impossible: a probabilistic approach is essential [*Newhall and Hoblitt*, 2002; *Sparks*, 2003; *Martin et al.*, 2004; *Marzocchi et al.*, 2007].

[3] The sizable uncertainty in scientific evidence makes the risk associated with volcanic activity difficult and stressful to manage. But the recent substantial progress in developing event tree and geostatistical methods for quantifying eruption probability [e.g., *Newhall and Hoblitt*, 2002; *Marzocchi et al.*, 2004, 2008; *Jaquet et al.*, 2006] allows a range of risk metrics to be calculated. These have the potential to rationalize decision making across a broad range of volcanological questions: (1) When should the call for evacuation be made? (2) What early preparations should be made for a volcano crisis? (3) Is it worthwhile to wait longer? (4) What areas should be covered by an emergency plan? (5) How should one handle the spatial uncertainty for calling an evacuation in calderas or volcanic fields?

[4] For some remote volcanoes, such questions may be answered conservatively, with no thought for risk assessment. For example, in a sparsely populated region, the emergency plan might extend for the entire geographical area affected by the largest historical eruption. However, for cities on volcanoes, sound answers do not come as readily as criticism. Resolution of difficult dilemmas is typically left to the wisdom and professional judgment of civil protection officials, who carry the burden of weighing fairly the advantages and disadvantages of any decision. In exercising this societal responsibility, appropriate volcano risk metrics (VRM) can assist civil protection officials to summarize, in a systematic, traceable and auditable manner, the key information relevant for balancing contrasting societal interests.

[5] VRM create a bridge between scientists who now have the techniques to estimate the probability of a threatening volcanic event [*Newhall and Hoblitt*, 2002; *Marzocchi et al.*, 2004, 2008; *Jaquet et al.*, 2006], and the decision makers who operate in the discrete world of Boolean logic: deciding either to take mitigating action or not [see, i.e., *Marzocchi and Woo*, 2007; *Woo*, 2008]. According to the pioneering Italian probability theorist, Bruno de Finetti: ''It is the evaluation of probabilities, unconscious or more or less conscious as it may be, that influences and determines everyone's behavior under uncertainty''. Here, we aim to

[1]Istituto Nazionale di Geofisica e Vulcanologia, Rome, Italy.
[2]Risk Management Solutions, London, UK.

Copyright 2009 by the American Geophysical Union.
0148-0227/09/2008JB005908$09.00

Figure 1. Part of the Campi Flegrei caldera (courtesy of G. Vilardo of INGV-Osservatoro Vesuviano). The area within the view, which is from the west, is approximately 30 km^2.

particularize such a statement in a volcanological context and to make the link between probability evaluation and decision making quantitative and transparent. Specifically, we present the basic VRM principles and apply them to address the five critical questions just mentioned. Such applications have the purpose of illustrating the practical usefulness of the VRM principles, which can be generalized to handle different potential problems such as, for example, deciding whether or not to build a new industrial facilities or lifelines in particular places [see, i.e., *Ho et al.*, 2006]. In section 2, we first introduce the VRM issue, and review some past efforts at tackling this problem. Then, we provide quantitative solutions to answer the questions, with illustrations from Mount Vesuvius and Campi Flegrei, representing two of the highest risk volcanic areas in the world (Figures 1 and 2; see also *Scandone et al.* [1993] for a pioneering work on long-term risk assessment at Mount Vesuvius). The examples aim to clarify practical VRM implementation; they represent just a first application of the procedures because many parameters used here may require substantial revision by Civil Protection, and the probability forecasting rules are presently being fine-tuned

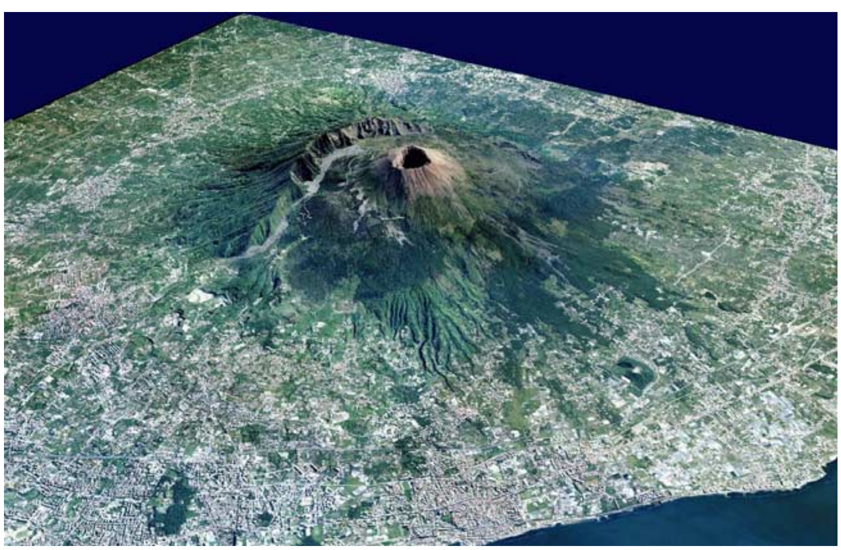

Figure 2. Mount Vesuvius and the surrounding area (courtesy of G. Vilardo of INGV-Osservatoro Vesuviano). The area within the view, which is from the southwest, is approximately 150 km^2.

Table 1. Glossary of the Most Relevant Terms and Acronyms Used in the Paper

	Definition
P	probability of a generic event
$\mathcal{V}$	the willingness to pay to save a human life
C	cost of an evacuation
L	the benefit in lives saved through an evacuation
E	proportion of people at risk who would owe their lives to the evacuation call
R	the average socioeconomic loss per capita in an evacuation
F	the fraction of the designated population killed in a volcanic eruption
$P^\star$	probability threshold for calling an evacuation
$P_{PF}^{(X)}$	the monthly probability that the cell X will be engulfed by a pyroclastic flow
P(VEI4+)	the annual probability of an eruption with VEI 4 or larger
Γ	the time at which the eruption probability reaches 90%
$t^\star$	the time of the call for evacuation
d	the time since the call for evacuation
$\mathcal{O}_m$	the relative odds that the threat at time T_m will decrease rather than decrease
$\mathcal{P}_A$	conditional probability of an eruption threatening the red zone (zone A)
$\mathcal{P}_B$	conditional probability of an eruption threatening the yellow zone (zone B)
$U(d)$	the proportion of the population in a designated region that remains unevacuated d days after an evacuation is called
$\mathcal{U}(T_m)$	the likelihood of unevacuation, if an evacuation commences at T_m
Π_τ	probability that a resident in the red zone around Vesuvius will be killed by a pyroclastic flow in a time window τ
VRM	volcanic risk metrics
GDP	gross domestic product
FN	frequency number
CBA	cost-benefit analysis

in an ongoing Italian project funded by Civil Protection (INGV/DPC V1 "UNREST"). In order to facilitate comprehension of this paper, a glossary of the most relevant symbols is reported in Table 1.

2. Basic Principles of Volcano Risk Metrics

2.1. Individual and Societal Risks

[6] Before going into details of VRM, it is worth discussing the recent past efforts at addressing volcanic risk, and the fundamental concept of individual risk versus societal risk. In reviewing VRM, a start is made with the risk to an individual from volcanic action. Individuals are exposed to all kinds of physical risks in their daily lives. These may be purely voluntary, as in playing dangerous sports, or they may be intrinsic in undertaking dangerous occupations. Individual risk criteria are established by authorities in circumstances where an unduly heavy risk is imposed on certain individuals. The rationale for individual risk criteria is clear: irrespective of the benefits to the national economy, it is grossly unfair to impose an excessive risk on specific workers such as underground miners. It is a mark of a modern risk-conscious society that no group of industrial workers should experience life expectancies far short of the societal norm.

[7] As professional experts, volcanologists are responsible for their own decisions to expose themselves to volcanic hazards. Such decisions are typically made without the support of any formal risk assessment. However, the deaths and injuries of volcanologists in tragic circumstances [e.g., *Williams*, 2001] indicate that even professionals are not so risk aware that they could not gain from an individual risk assessment. Certainly, the general public is in need of guidance on risk exposure. In particular, volcano tourists should be warned if, on a particular day, the accident risk is significantly greater they regularly experience. Indeed, depending on the individual risk level, the higher slopes of an active volcano may be closed to visitors for a while during a period of crater activity.

[8] For a densely populated region around a volcano, permanent inhabitants may greatly outnumber temporary visitors. Compared with the closure of a volcano to tourists, any decision on population evacuation would be much more fraught, since it should involve careful consideration of societal risk issues. The loss of a single life is a family tragedy; the accidental loss of many lives may be a societal disaster. VRM have been devised to address the concern over multiple fatalities.

[9] A societal risk criterion, which has been quite widely adopted, is based on the expression of risk in terms of the frequency and severity distribution of the number of fatalities. Let $F(n)$ be the frequency of accidents with n or more fatalities. Then a plot of $F(n)$ against n is called an FN curve. Typically, because of the rarity of great disasters killing many people, this curve is plotted on a log-log scale. FN curves are a convenient visual means of depicting fatality distributions, and have also been used to define safety tolerability criteria. In a volcanological context, FN curves can be computed using probabilistic volcano risk assessment methods, and such curves were first developed by *Aspinall* for application to the ongoing volcanic crisis at Montserrat [*Aspinall et al.*, 2002; see also *Baxter et al.*, 2008].

[10] The simplest FN risk criterion is defined by a diagonal line of negative slope on an FN plot. The region above the line is indicative of excessive risk, which might be deemed to be societally intolerable. Points in this dangerous region correspond to hazardous situations where there are far too many events killing a few; or many events killing a moderate number, or a significant number of events killing a large number of people. By contrast, the region below the line is indicative of moderate risk, which might be deemed to be societally tolerable.

[11] While an FN criterion has been used for some safety regulations, there is no clear systematic empirical or quantitative basis for the precise choice of this criterion. Furthermore, there are some conceptual difficulties which detract from the use of an FN criterion as a risk metric [*Evans and*

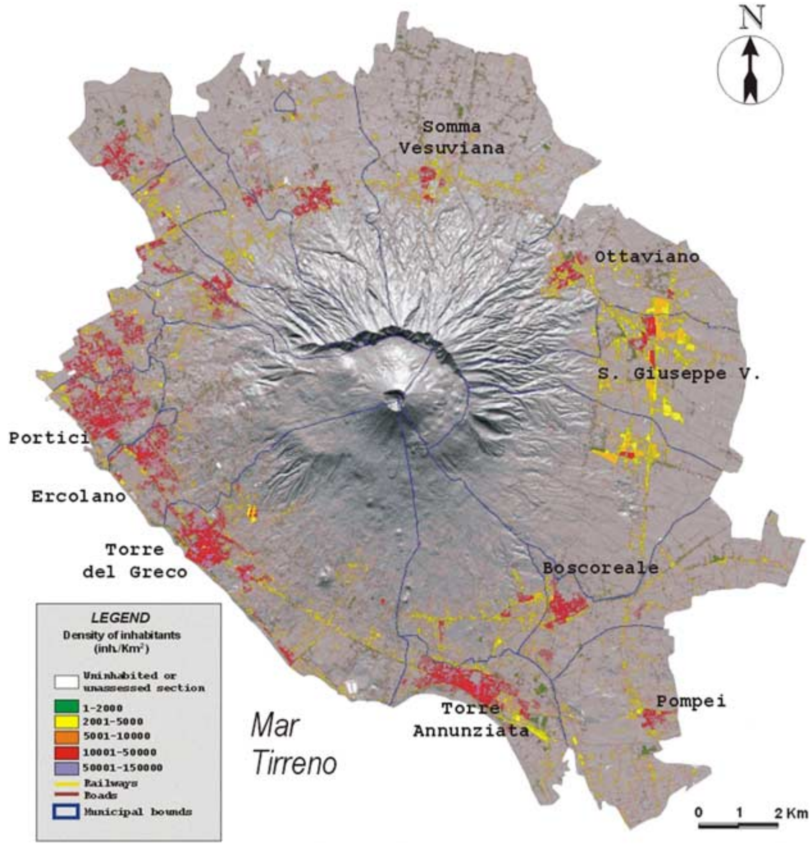

Figure 3. Red zone of Mount Vesuvius showing the density of inhabitants [see *Pareschi et al.*, 2000].

Verlander, 1997]. Specifically, *Evans* [2003] has pointed out that consistent decision making under uncertainty requires that the risk criterion should be the statistically expected value of some function of the number of fatalities. This may seem an arcane point of risk theory warning of potential anomalies, but it points toward cost-benefit analysis of risk mitigation action, where the expected reduction in the number of fatalities is weighed against the expected economic dislocation costs. Indeed, as *Fischoff et al.* [1981] have observed in their treatise on acceptable risk, risk is only acceptable if some benefit can compensate for the risk.

2.2. Cost-Benefit Analysis

[12] The trade-offs involved in taking mitigating action in the interests of public safety can be studied within the economic framework of cost-benefit analysis (CBA). Consider a situation where a decision maker has to choose between two actions: either (1) protect or (2) do not protect. The cost of protection is C. In the absence of protection, the decision maker incurs a loss L, which exceeds C, if an adverse hazard state arises. Let the probability of the adverse hazard state arising, within a specified time window, be denoted by P. If the expected expense is to be minimized, then the optimal policy is

$$\begin{aligned} \text{To protect, if} \quad & P > C/L, \\ \text{Not to protect, if} \quad & P < C/L \end{aligned} \tag{1}$$

The minimal expense is then $\min\{C, PL\}$. In a volcanological context, protection may involve evacuation, which carries a cost of C. The adverse hazard state here is one of volcanic eruptivity, for which a decision not to protect carries a large loss penalty of L, measured in human fatalities.

[13] *Viscusi* [1992] has reviewed situations where the avoidance of fatalities is a sensitive public policy issue. As we shall see in this paper, there are different quantitative formulations of this principle, according to the problem under discussion. A quantitative CBA requires the evaluation of both costs and benefits in the same common unit, the most convenient being monetary value. In some cases this is straightforward, but in other cases this evaluation is far from obvious. A standard practical approach [*Viscusi*, 1992] is to determine the amount of money authorities are willing to pay to save a human life. There is natural public antipathy to the notion of placing a monetary value on human life, but this is indirectly measurable by the willingness to pay to avoid danger.

[14] As an example, consider the important mitigation action taken by authorities to reduce the long-term risk at Mount Vesuvius. The name of the initiative is VesuVia (literally meaning "away from Vesuvio") and it essentially consists of offering 30,000 Euro to a family to move away from red zone, which is the designated area for evacuation before an eruption (Figure 3). In this case, the implicit monetary value of a life saved is the price paid per person to mitigate the risk, divided by the probability of being killed by the hazard. Let us assume that the average number of people in a family is 3.5; the offer is then tantamount to about 8600 Euro per person. The likelihood of a long-term red zone resident dying in a Vesuvius eruption is of the

order of 0.01, based on the historical experience of the past several millennia (see Appendix A). Thus the willingness to pay to save a human life ($\mathcal{V}$ hereafter) in this Italian context is the ratio between the money paid (8,600 Euro) and the probability of the deadly event (0.01), i.e., about 860,000 Euro. Allowing for an uncertainty margin in the underlying assumptions of threat level and future life expectancy, a minimum value of $\mathcal{V}$ might reasonably be taken to be 800,000 Euro. This minimum is consistent with values used for risk studies in affluent industrialized nations.

[15] From a public policy perspective of fairness, the minimum value of $\mathcal{V}$ may be taken to be the same for all individuals. However, the false alarm cost C is different for visitors and local residents. Visitors have a home and livelihood elsewhere, and may return with comparatively little dislocation to their daily lives. By comparison, the evacuation cost for local residents is much more burdensome. Even if temporary alternative shelter can be obtained, their livelihoods may be severely disrupted for a number of months. Scientists are in a special category, since they may incur additional professional costs if precluded from working on an active volcano. False alarms have also intangible costs, mostly related to a possible loss of credibility in scientists, because society may overestimate their predictive capability. We argue that this critical problem can be significantly tackled with an extensive educational program.

[16] From a decision maker's perspective, CBA can be seen as a framework to define some key probability thresholds. A first exposition of cost-benefit analysis (CBA) in the context of the call for volcano evacuation has been given by *Marzocchi and Woo* [2007] and *Woo* [2008]. But this is by no means the only difficult decision which cost-benefit analysis can rationalize. At lower levels of threat, CBA can justify preparatory measures in advance of a crisis. These might include issues such as transportation logistics and emergency service preparedness. Furthermore, just as importantly, CBA can be used even outside crisis periods to justify and refine the geography of emergency planning.

3. When Should the Call for Evacuation Be Made?

[17] The recent simulation exercises of volcanic eruption, Major Emergency Simulation Exercise (MESIMEX) and RUAUMOKO (http://www.exerciseruaumoko.co.nz/), close to the large cities of Naples and Auckland have affirmed the understandable tendency for some volcanologists to call for an evacuation when the probability of events is reasonably high [*Marzocchi and Woo*, 2007; J. Lindsay et al., Real-time eruption forecasting in the Auckland Volcanic Field: application of BET_EF during the New Zealand National Disaster Exercise "Ruaumoko", submitted to *Bulletin Volcanology*, 2008]. For a volcanologist, caution may be an optimal strategy for himself or herself: no scientist wants to be accused of acting without full evidence, and causing a false alarm. And should an eruption occur, the same defense can be made as after the 2004 Indian Ocean tsunami: scientists are trained in determinism, and want to be sure before expressing a public opinion [*Woo and Aspinall*, 2005].

[18] However, this attitude is far from optimal for those at risk, as is appreciated by experienced volcanologists familiar with the literature on acceptable risk (e.g., C. G. Newhall, personal communication, 2008). As an example, suppose that the probability of occurrence of a pyroclastic flow is 5%. Prima facie, this is not sufficient for a volcanologist to call for an evacuation. However, looking at the same problem from the perspective of an inhabitant at risk, 5% may well seem to be an intolerably high probability. Suppose further that the probability of occurrence of a pyroclastic flow is 10%. This is still a low level of confidence for a scientist, but as demonstrated by *Marzocchi and Woo* [2007] and *Woo* [2008], the corresponding societal risk may well be high enough to warrant an evacuation, according to CBA.

[19] In order to illustrate the problem in practice, we summarize the case of Vesuvius. *Woo* [2008] shows that the ratio C/L can be written as

$$\frac{C}{L} = \frac{NR}{EN\mathcal{V}} = \frac{(R/\mathcal{V})}{E} \tag{2}$$

where R is the average socioeconomic loss per capita, N is the number of people involved in the evacuation, and E is the proportion of people at risk who would owe their lives to the evacuation call, should the adverse eruptive hazard state materialize. From section 2.2, we estimate a reasonable minimum value for $\mathcal{V}$ is 800,000 Euro. This is approximately 25 times GDP per capita, or 100 times GDP per capita for a 3-month evacuation period. GDP is the Gross Domestic Product as published in World Bank reports, and it is an appropriate economic indicator for socioeconomic loss expressed in terms of US currency. Specifically, for a period of evacuation of three months in Italy, R is about 8,000 Euro and $R/\mathcal{V}$ is about 0.01. In practice, equation (2) can be written as

$$\frac{C}{L} = \frac{0.01}{E} \tag{3}$$

[20] From Vesuvius pyroclastic flow simulations [*Neri et al.*, 2007] (see Appendix A), the percentage of area surrounding Vesuvius that will be impacted by one pyroclastic flow is estimated at between 10% and 20%. The parameter E may be a function of the time elapsed since the evacuation is called (see section 4S). Here, as an example, we take $E = 0.10$. This leads to

$$\frac{C}{L} \equiv P^{\star} = 0.1 \tag{4}$$

which is the probability threshold to be overcome ($P^{\star}$) for calling an evacuation [see *Marzocchi and Woo*, 2007]. The underlying principle is simple: the higher the willingness to pay to mitigate the risk suffered by an endangered population ($\mathcal{V}$), the lower should be the threshold for triggering an evacuation $P^{\star}$. Similarly, the more people whose lives would depend on a successful evacuation call (higher E), the lower should be this threshold. It should be noted that the provision of economic incentives to evacuate, such as giving out food rations or a living allowance, effectively reduces the economic cost to the evacuees, and hence lowers the threshold for evacuation. Thus, in some developing countries where such practical incentives are given, the probability threshold for evacuation may be lower than in the industrialized world.

[21] To summarize, high levels of scientific confidence in the imminence of an eruption should not be a prerequisite for an evacuation call. Precautionary decision making may require mitigating action to be taken in the face of considerable scientific uncertainty. The contribution of scientists must be judged in this risk context. A posteriori criticism of volcanologists in the wake of a false alarm should recognize their efforts to serve not their own, but the broader interests of society, in particular to safeguard human life as far as possible. An illustrative example is Tungurahua, Ecuador, in 1999, where the volcanologists involved in that crisis preferred to err on the side of giving an early alert leading to a temporary evacuation, even at some cost to credibility.

4. What Early Preparations Should Be Made for a Volcano Crisis? Different Stages for a Progressive Prioritized Evacuation

[22] Even in the early stages of a volcanic crisis, while the eruption probability P is still quite low, actions may be taken preparatory to a possible major escalation of the crisis. CBA can provide guidance as to which actions, short of a population warning, might be warranted at a particular moment. For example, emergency services may be placed on high alert, auxiliary public transportation may be commandeered, and some safety-critical engineering systems may be shut down. In regard to the latter, there is a parallel with seismic safety, where CBA has been used in California to control dam reservoirs as a precaution against the shaking of a sudden large earthquake. Similarly, at low P values, a range of low-cost preparedness measures would be justified by CBA, as a precaution against volcanic action.

[23] In addressing the problem of calling an evacuation, one basic component is all too commonly neglected: the time taken for an evacuation. Usually, the evacuation process is considered a deterministic process that can be expressed in terms of the number of people evacuated in a fixed time period. Past experience has shown that such a process is far from being deterministic; one that can be better described by a probability distribution with a fat tail. Amongst the human stragglers late to evacuate are those with physical or mental frailties, those hesitant to leave for fear of looting, and those sceptical of scientific advice, preferring instead to put their faith in religion. A recent example of such human behavior is the ill-fated evacuation from Hurricane Katrina of New Orleans, which is a city with some socioeconomic similarities with Naples. Let $U(d)$ be the proportion of the population in a designated region that remains unevacuated d days after an evacuation is called. A transportation computer simulation will generate a residual nonevacuation time curve of the form shown in Figure 4. Consider the way in which the state of a volcano evolves toward an eruption (see Figure 5). The possible evolutions of this process define a class of trajectories over time of the eruption probability P per time window. These may be termed p trajectories. Some possible discrete trajectories, varying in the rapidity in reaching unit probability from the present state, are illustrated in Figure 5 and are described here. Here, we assume that each eruption (or deadly event) occurs when P $\geq$ 90%. We define $t^\star$ as the time of the call for evacuation, Γ is the time at which P reaches 90%, and $d^\star$ is the lead time $\Gamma - t^\star$. The plot shows three possible patterns to eruption. Case 1 has a gradual increase of P with time; case 2 has some major steps when the monitoring system detects the presence of magma responsible for the unrest, and when magma starts to move toward the surface; case 3 mimics a rapid P escalation few days before the event. Case 3 is typical for open conduit volcanoes [*Marzocchi and Zaccarelli*, 2006] where precursory activity may be very close in time to the eruption; case 1 and case 2 are more typical for closed conduit volcanoes (ibid).

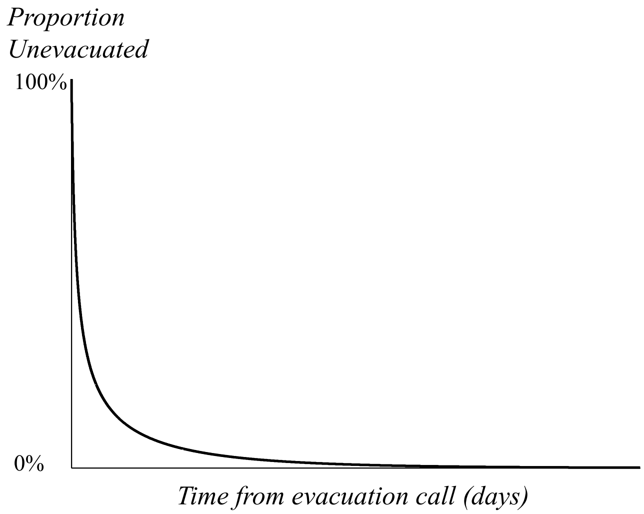

Figure 4. This is an illustrative plot charting the decline of the percentage of unevacuated people as a function of the time since the call of an evacuation. No specific scale is indicated on this generic plot because this varies from one volcano to another.

[24] If we assume a probability threshold of $P^\star$ for calling an evacuation, each p trajectory attains this threshold at different times, therefore the lead times are: $d_1^\star = \Gamma - t_1^\star$, $d_2^\star = \Gamma - t_2^\star$, and $d_3^\star = \Gamma - t_3^\star$. Therefore, for the kth p trajectory (k = 1, 2, 3), $d_k^\star$ is the number of days within which there is evident notice of an imminent eruption. Figure 5 illustrates the definition of $P^\star$, $d_k^\star$, and $t_k^\star$. Let ω_k be the relative likelihood of the kth trajectory. Theoretical geophysical models of the stochastic eruption process, e.g., probabilistic time-to-failure models (see *Voight and Cornelius* [1991] and *Kilburn* [2003] for theoretical background), as well as empirical observations of past events can be used to estimate ω_k. Then, the expected proportion of the population unevacuated after d days since the evacuation call is

$$\langle U(d) \rangle = \sum_k \omega_k U(d_k) \tag{5}$$

[25] The expected proportion of the designated population N who would owe their lives to an evacuation is

$$N[1 - \langle U(d) \rangle]F \tag{6}$$

where F is the fraction of the designated population engulfed by a pyroclastic flow, or otherwise killed in a volcanic

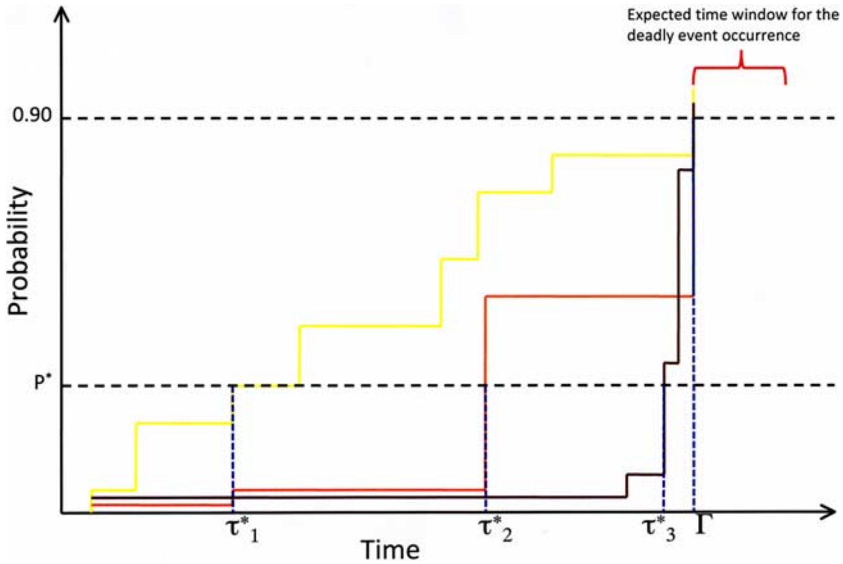

Figure 5. Eruption probability evolution as a function of time for three possible preeruptive patterns (see text for more details): case 1 (yellow line), case 2 (orange line), and case 3 (dark red line). $P^\star$ is the probability threshold for evacuation defined through CBA. The instants $\tau_1^\star$, $\tau_2^\star$, and $\tau_3^\star$ represent the evacuation time for case 1, case 2, and case 3, respectively.

eruption by roof collapse, ballistic ejecta, asphyxiation, heart failure, etc. The parameter E of equation (2) is the asymptotic value (for $d \to \infty$) of

$$E(d) = [1 - \langle U(d) \rangle]F \tag{7}$$

[26] A reasonable value for F is about 0.1–0.2 (see Appendix A). Allowing for the realistic prospect that the unevacuated proportion $\langle U \rangle$ may be relatively significant, lying between 0.1 and 0.2, the expected proportion of the red zone population who would owe their lives to an evacuation E is approximately 10%. In a situation where a significant proportion of the population were to resist evacuation, then the probability threshold for calling an evacuation would be raised. A practical feature of CBA is its capability to account for such aspects of population behavior. Note that the designated population can be segmented by geographical location around the volcano (e.g., red zone, yellow zone, etc.), and also by socioeconomic factors. Such segmentation allows a CBA to optimize a strategy for progressive prioritized evacuation, or non-evacuation. For example, a higher proportion of frail people will be liable to be killed in an eruption from ancillary factors such as respiratory and heart ailments. Accordingly, the probability threshold obtained by CBA for their evacuation will be lower, and it would generally be prudent for them to be evacuated early.

[27] For example, let us suppose that human beings can be categorized into two broad groups depending on their sensitivity to critical conditions like respiratory difficulties. As a simple categorization, we define the group A as composed of people with robust health, and group B as composed of more frail people. Because F may be significantly higher for the latter group, E would be higher (see equation (7)), and the probability threshold for evacuation of group B would be lower (see equations (3) and (4)).

5. Is It Worthwhile Waiting for the Next Volcanological Alert Bulletin?

[28] At various times during a volcanic crisis, volcanological bulletins will be issued. Consider the situation at time T_m during a volcanic crisis, when the mth volcanological alert bulletin has been issued. If the imminent eruption probability P_m is below the threshold for evacuation, then no action need be taken until the next bulletin at time T_{m+1}.

[29] But, suppose that P_m exceeds the threshold for evacuation. Rather than evacuate straight away, civil protection authorities may be inclined to wait for the next bulletin at time T_{m+1}, to see if the threat is actually rising, rather than just having peaked. How should this decision to delay be made rationally?

[30] Let the threat level at time T_m be designated by the eruption probability P_m. As before, let $U(d_m)$ be the fraction of unevacuated people if an evacuation is authorized and commenced at T_m, and d_m is the time left before the event occurs. This quantity can be set to a probability function

$$U(d_m) \equiv \mathcal{U}(T_m) \tag{8}$$

where $\mathcal{U}(T_m)$ represents the likelihood of unevacuation, if an evacuation is authorized and commenced at T_m. The probability function $\mathcal{U}(T_m)$ will vary, as well as $U(d_m)$, according to population group, and may be significantly high for the frail.

[31] The expected cost of impatience in not waiting for the next bulletin is

$$\text{Prob}[P_{m+1} < P_m]C \tag{9}$$

This is the probability that the threat is diminishing, multiplied by the cost of an unnecessary evacuation. On the other hand, the expected benefit of not delaying to the next bulletin is

$$\mathrm{Prob}[P_{m+1} > P_m][\mathcal{U}(T_{m+1}) - \mathcal{U}(T_m)]L \tag{10}$$

In equations (9) and (10), C and L have the same meaning as in equation (2). This is the probability that the threat is rising, multiplied by the marginal likelihood of unevacuation (given the rising threat), multiplied by the value of lives saved.

[32] Let $\mathcal{O}_m$ be the relative odds that the threat at time T_m will decrease rather than decrease at time T_{m+1}, i.e.,

$$\mathcal{O}_m = \frac{\mathrm{Prob}[P_{m+1} < P_m]}{\mathrm{Prob}[P_{m+1} > Pm]} \tag{11}$$

Then the evacuation should proceed straight away, provided the benefit (equation (10)) outweighs the cost (equation (9)), i.e.,

$$[\mathcal{U}(T_{m+1}) - \mathcal{U}(T_m)] > \mathcal{O}_m C/L \tag{12}$$

As before, we can set an upper bound of 0.1 for C/L. Taking into account equations (9)–(12), it is not worthwhile delaying the evacuation if

$$[\mathcal{U}(T_{m+1}) - \mathcal{U}(T_m)] > \mathcal{O}_m 0.1 \tag{13}$$

An assessment of $\mathcal{O}_m$ would depend on the crisis history up to time T_m. A time series analysis may be used to estimate $\mathcal{O}_m$. If the threat has been steadily rising, then $\mathcal{O}_m$ will be small. However, if the threat has been fluctuating, then $\mathcal{O}_m$ may be evens. Suppose the odds of the threat decreasing are no better than evens (i.e., if the outlook is not optimistic). Then it would be worth evacuating now if

$$[\mathcal{U}(T_{m+1}) - \mathcal{U}(T_m)] > 0.1 \tag{14}$$

Where effective evacuation logistics are in place, so that the marginal likelihood of unevacuation is small, then a delay until the next bulletin might well be warranted. But where, as around Naples, traffic might become gridlocked by road blockages, a delay of a day or two might have a significant impact on evacuability. Thus, evacuation would typically be advised when the evacuation criterion is attained. This is especially the case for vulnerable groups, e.g., the sick and the elderly, at greater risk of harm from an eruption.

6. What Areas Should Be Covered by an Emergency Plan? The Geography of Evacuation Decision Making

[33] One of the challenges for a highly urbanized area is to define the spatial extent of the evacuation plan. For the case of Vesuvius, Italian Civil Protection, assisted by a panel of expert volcanologists, calibrated the emergency plan using a reference scenario based on a sub-Plinian event (VEI 4), similar to the most recent eruptions which have ended a closed conduit period like the present one [*Rosi et al.*, 1993]. However, it has been argued by some volcanologists [e.g., *Mastrolorenzo et al.*, 2006] that the Emergency Plan should account also for the occurrence of the so-called worst expected eruption, namely the Pompeii (79 A.D.) or the Avellino (3780 B.P.) eruption.

[34] This open debate [*Hall*, 2007] over the evacuation geography, which appears scientifically inconclusive, can be resolved rationally and concisely using CBA. Consider a generic volcanic crisis situation, where there is a red zone, proximal to the volcano, which is threatened by a major eruption, designated as scenario A, and an outer yellow zone, which is threatened by a rarer great eruption, designated as scenario B. In our simplified case, we consider only these two scenarios with conditional probability $\mathcal{P}_A$ and $\mathcal{P}_B$, respectively, and $\mathcal{P}_A + \mathcal{P}_B = 1$. Note that the term worst is avoided, because it is not scientifically possible to put an absolute reasonable upper bound on the size of an eruption. Instead, we need to consider a set of scenarios (two in our case), each having a conditional probability of occurrence.

[35] Let the CBA probability threshold for evacuating the inner red zone be denoted as $P^\star$ (A), and let the corresponding CBA probability threshold for evacuating the outer yellow zone be denoted as $P^\star$ (B). The proportion of the population in the yellow zone who would owe their lives to an evacuation is expected to be less than for the red zone for several reasons:

[36] 1. Proportionately fewer would be in harm's way from the path of a directed pyroclastic flow, because of the attenuation of its intensity with distance. Thus the F factor is lower.

[37] 2. Proportionately fewer might be expected to sense the danger and heed instructions to evacuate. Thus an evacuation call would be less effective in its response, and the $\langle U(d)\rangle$ factor is higher.

[38] In practice, this means that

$$P^\star(B) > P^\star(A) \tag{15}$$

Let the probability of an eruption of size A or higher be P. Then the probability of a great eruption is $P\mathcal{P}_B$. Whereas for an evacuation of the red zone to be justified on a cost-benefit basis, P would have to exceed the threshold $P^\star(A)$, for an evacuation of the yellow zone to be justified on a cost-benefit basis, $P\mathcal{P}_B$ would have to exceed $P^\star(B)$. Clearly, if it turns out that $P^\star(B) > \mathcal{P}_B$, then this broader evacuation would never be justified. Given the relative rarity of a great eruption, there could never be enough confidence in the imminence of such a great event to justify an evacuation of the yellow zone.

[39] To illustrate this prospect, consider Vesuvius. For the red zone, as assessed in section 3, we have $P^\star(A) = 0.1$ [*Marzocchi and Woo*, 2007]. Since the conditional probability of a larger Plinian eruption (VEI 5+) $\mathcal{P}_B$ is about 0.1 [*Marzocchi et al.*, 2004], $P^\star(B)$ would have to be less than 0.1, for an evacuation of the yellow zone ever to be warranted. This would only happen if more than 10% of the population in the yellow zone would owe their lives to an evacuation call (see equation (2)). This is assessable using pyroclastic flow models, corroborated by the geological evidence of the pyroclastic flow footprint from the Avellino eruption [*Mastrolorenzo et al.*, 2006]. With the current state of information, such a high proportion as this would seem rather unlikely. Thus CBA reasoning supports the geogra-

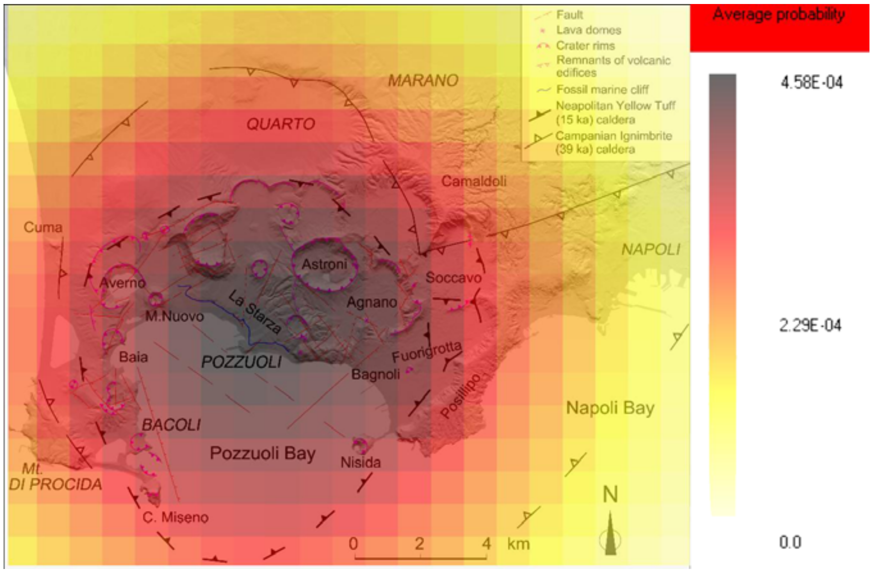

Figure 6. Probability map for pyroclastic flow occurrence during a quiet period of Campi Flegrei caldera. Probabilities are per month. Details of calculation are reported in Appendix B.

phical restriction of the existing civil protection emergency plan to the red zone.

7. How Should Evacuation in Calderas and Volcanic Fields Be Managed?

[40] In the case of calderas or volcanic fields like Campi Flegrei (Italy), Long Valley (United States), and the Auckland Volcanic Field (New Zealand), it is not possible to define a priori a fixed geometry of the evacuation area due to the additional uncertainty over the next vent location. Fortunately, the principles expounded in this paper can be used also to address this more complicated problem.

[41] Consider a generic caldera having a large geographical extent. At any given time during a crisis, suppose that volcanologists compute the probability that a threatening event (for instance, the occurrence of a pyroclastic flow) will emanate from a specific location of the caldera. Because of the size of the caldera and the range of possible vents, there is no predefined evacuation zone, but the geography of such a zone can be defined through CBA.

[42] Let us subdivide the caldera into a grid of small cells. Consider a generic cell **X** of this grid. The issue is whether this cell should be included in the overall evacuation zone. This question is resolved in the following way. Let us assume that the threatening event is the arrival of a pyroclastic flow in this cell **X**. The probability of such an event $P_{PF}^{(\mathbf{X})}$ can be calculated a Bayesian event tree structure [*Newhall and Hoblitt*, 2002; *Marzocchi et al.*, 2004, 2008], where the probability $P_{PF}^{(\mathbf{X})}$ can be disaggregated as a product of seven component probabilities.

$$P_{PF}^{(\mathbf{X})} = p_1 p_2 p_3 \sum_k \sum_S \left[p_4^{(k)}\ p_5^{(S)} p_6^{(S)} p_7^{(k,S,\mathbf{X})} \right] \quad (16)$$

where p_1 is the probability of unrest in the next month; p_2 is the probability that the unrest is due to magma reactivation, given an unrest is detected; p_3 is the probability that the magma reaches the surface, given magmatic unrest; $p_4^{(k)}$ is the probability of having that eruption in the *k*th cell, provided that there is an eruption; $p_5^{(S)}$ is the probability of an eruption with size *S*, given an eruption occurred in the *k*th cell; $p_6^{(S)}$ is the probability of occurrence of a pyroclastic flow, given an eruption in the *k*th cell with size *S*; $p_7^{(k,S,\mathbf{X})}$ is the probability that cell **X** will be struck by a pyroclastic flow originating from an eruption of size *S* in the *k*th cell. If cell **X** is sufficiently small, $E_{\mathbf{X}}$ can be considered about 1, i.e., all people in the cell will owe their lives to an evacuation. In this case, equations (2)–(4) can be rewritten as

$$P_{PF}^{\star(\mathbf{X})} = \frac{C}{L} = \frac{R/\mathcal{V}}{E_{\mathbf{X}}} = 0.01 \quad (17)$$

In practice, the cell **X** should be included within the evacuation zone if $P_{PF}^{(\mathbf{X})}$ exceeds 0.01. The overall evacuation zone is then the set of all cells **X** for which this inequality holds. In the limiting case where there is only a single vent, this analysis yields the criterion given in section 6.

[43] The presence of many possible vents provides quantitative justification for authorities waiting longer before taking any action, especially where the vents have a sizable separation distance. Because the probability of an eruption is spatially diffused among a number of possible vents, the evacuation criteria will generally take longer to be exceeded. In the limiting case where there is only one caldera vent, this supports the distinction between red and yellow zones for emergency planning, and could be used for further refining the zone boundary.

[44] This CBA allows the extensive volcanological hazard mapping in progress on Campi Flegrei to be maximally utilized for Civil Protection purposes. It is worth stressing that this procedure can be established during the current period of comparative quiescence, well in advance of the next crisis. For illustrative purposes, (with further details given in Appendix B), suppose that the evacuation area is defined

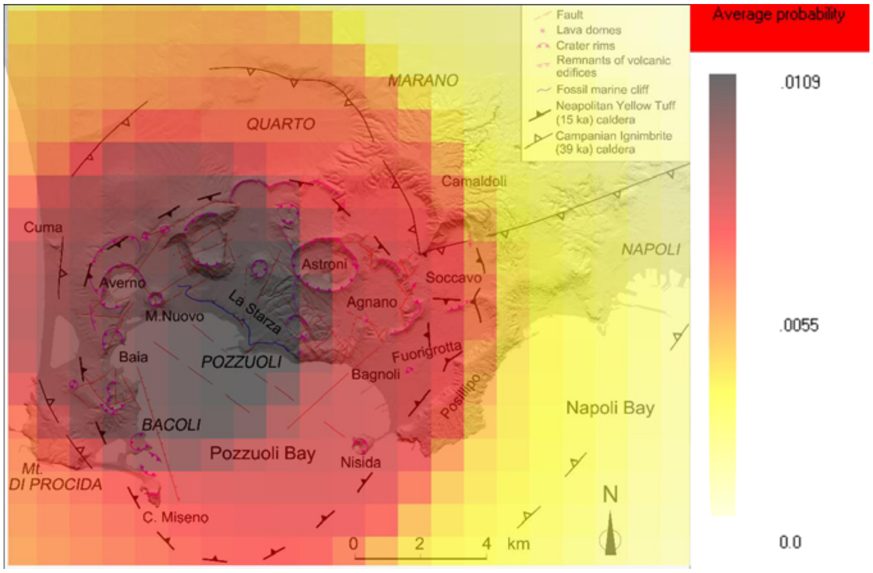

Figure 7. The same as Figure 6 but relative to an episode of unrest with monitoring anomalies observed below Monte Nuovo. Note the different probability scale compared to Figure 6.

only by considering the pyroclastic flow hazard. In reality, ash fall and lahars should also be taken into account. A preliminary analysis has been undertaken based on empirical models and observations. The map on Figure 6 shows the probability per month for each spatial cell to be reached by a pyroclastic flow. The estimation is made during a quiet period of the caldera. Let us then suppose that an episode of unrest takes place with most monitoring observations indicating the vent location to be close to Monte Nuovo. This modifies the vent opening map and, consequently, the probability that a generic cell will be reached by a pyroclastic flow. Figure 7 is a revised version of Figure 6, taking into account the monitoring observations made during the episode of unrest. Figure 8 shows the same map as Figure 7, identifying the cells where the threshold of 0.01 for $P_{PF}^{\star(\mathbf{X})}$ is exceeded. Note that the evacuation area is different from what might have been decided from Figure 6, i.e., just from looking at the past activity of the caldera.

8. Conclusions

[45] The use of volcano risk metrics (VRM) to assist decision makers opens up a new era in quantitative volcano risk management. Through the VRM methods explained here, the geography of evacuation zones, both for a volcano and a caldera, can be delineated; the call for evacuation can

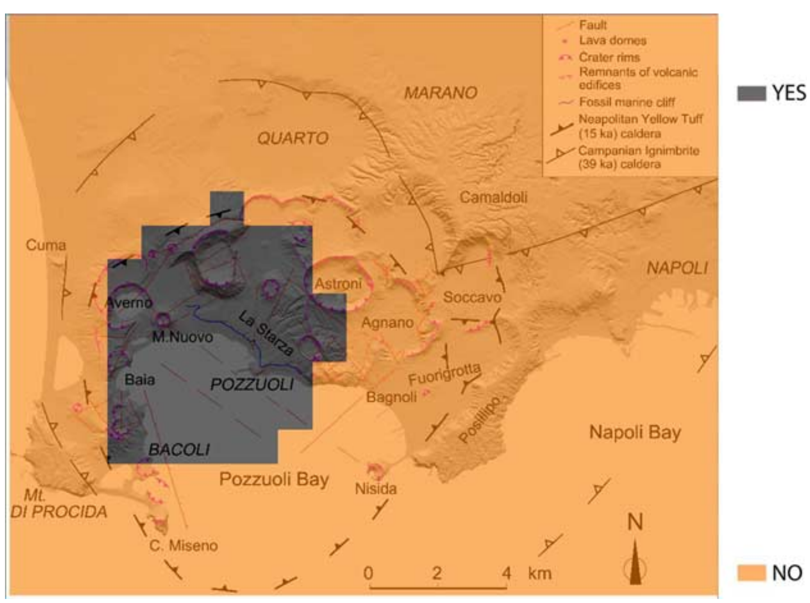

Figure 8. Map that relates to the same situation as Figure 7, but cells are distinguished according to their probability below or above the threshold for evacuation. In particular, all cells with a probability larger than 0.01 are in gray and represent the area to be evacuated.

be timed; and preparatory actions early in a volcanic crisis can be prioritized. VRM make full utilization of observational data and computer modeling analysis. As a result, volcanological monitoring and academic research are brought to the heart of the decision-making process. In particular, VRM creates a bridge among scientists who estimate probabilities of occurrence, and decision makers who set rules for cost/benefit analysis (CBA), which may well reflect some extraneous intangible political and civil defence factors. With a basic level of the willingness to pay to save a human life, the adoption of CBA is expected to save lives but not at an unreasonable economic cost. CBA provides a rational basis for calling for an evacuation before volcanologists could be really sure an eruption was going to occur, without resorting to qualitative arguments.

[46] The adoption of a systematic decision-support procedure has the advantage of providing a transparent audit trail, which reduces the degree of subjective opinion in volcanological communication with civil authorities. Should an evacuation turn out to be a false alarm, this audit trail would provide a direct means for tracking the decision process in any subsequent formal enquiry. The general public needs to understand that false alarms are sounded because their lives are highly valued. Conveying this message effectively will require a public education plan. Last, but not least, a formal quantitative scheme makes possible a program of future improvements not available to qualitative subjective decision making.

Appendix A: Likelihood of a Red Zone Resident Being Killed by a Pyroclastic Flow

[47] Here we estimate the probability Π_τ that a resident in the red zone around Vesuvius will be killed by a pyroclastic flow in a time window τ. At first, we set $\tau = 40$ years; this time window can be considered the average remaining life for an inhabitant (half of the average Italian lifespan). Next, in accordance with past activity of Mount Vesuvius, we assume that a pyroclastic flow can be produced only by VEI 4+ eruptions.

[48] The probability Π_{40} is then calculated as

$$\Pi_{40} = P(\mathrm{VEI4+})40F \tag{A1}$$

where P(VEI4+) is the annual probability of an eruption with VEI 4 or larger, and F is the portion of the area surrounding Vesuvius that will be hit by the pyroclastic flow. This is a reasonable approximation of Π_{40}, since the forecasting time window (40 years) is much less than the typical recurrence time of VEI 4+ events (about 700 years).

[49] $P(VEI4+)$ can be estimated by the frequency of past similar events; three VEI 4+ eruptions occurred in the last 2000 years (79 A.D., 472 A.D., 1631 [see *Scandone et al.*, 1993]); therefore $P(VEI4+) \approx 0.0015$ a^{-1}. F can be estimated visually by scanning the numerical simulations carried out by *Neri et al.* [2007] that show about 10–20% of the surrounding area of Mount Vesuvius is liable to be engulfed by a pyroclastic flow. Here we set $F = 0.15$.

[50] Then, equation (A1) becomes

$$\Pi_{40} = 0.0015 \times 40 \times 0.15 \approx 0.01 \tag{A2}$$

This value is only indicative because of the significant uncertainties. However, it represents a reasonable estimate for the purposes of this paper.

Appendix B: Example of Evacuation at Campi Flegrei During an Episode of Unrest

[51] The Campi Flegrei caldera is divided into a grid with cells 1 × 1 km. For a generic ith cell, the monthly probability $P_{PF}^{(i)}$ that the cell will be engulfed by a pyroclastic flow is calculated by using the event tree scheme reported by *Marzocchi et al.* [2004, 2008], and in equation (16) in the text. (Probabilities from *Marzocchi et al.* [2004, 2008] are considered as distributions to account for uncertainties. Here, since these calculations are illustrative, we consider best estimate probabilities). In this example, we consider two states of the caldera: a quiet period (period I), and an episode of unrest (period II).

[52] For period I, we set $p_1 = 8 \times 10^{-3}$ month^{-1}, $p_2 = 0.2$ (most episodes of unrest are due to instabilities of hydrothermal system rather than magma movement), $p_3 = 0.5$, $p_4^{(k)}$ is assigned to account for the tectonic setting of the caldera and the location of the past 21 eruptions, $p_5^{(S)}$ is 0.14, 0.52, 0.28, 0.06, respectively, for eruptions with VEI 2-, VEI 3, VEI 4, VEI 5+ (determined from past event sizes), $p_6^{(S)}$ is 0.05, 0.35, 0.75, 0.75 for the above four sizes [*Newhall and Hoblitt*, 2002], $p_7^{(k,S,i)}$ is assigned by using the empirical model dispersion provided by *Newhall and Hoblitt* [2002].

[53] For period II, the product $p_1p_2tp_3$, i.e., the probability of eruption is set to 0.2. $p_4^{(k)}$ is modified according to *Marzocchi et al.* [2008] to take into account the observation that all monitored anomalies are detected below Monte Nuovo. All other probabilities remain the same. (Notably, at present, the monitored parameters do not provide updated information as to the size of the impending eruption [*Sandri et al.*, 2004]). We should stress that these are purely indicative values, emerging from preliminary analyses carried out in the framework of a past INGV/DPC project (V3_2 Flegrei) and presently under revision in the new INGV/DPC V1 UNREST project both funded by Italian Civil Protection.

[54] Figure 6 shows the corresponding probability of occurrence per month of a pyroclastic flow at Campi Flegrei. Figure 7 is the same map calculated during a hypothetical episode of unrest with almost all anomalies recorded below Monte Nuovo. In this case, p_1 becomes 1 (definite unrest), and the $p_4^{(k)}$ values account for the location of monitored anomalies according to the scheme proposed by *Marzocchi et al.* [2008]. Figure 8 highlights the cells shown in Figure 7 where $P_{PF}^{(i)} > 0.01$, which is the designated threshold for evacuation.

[55] **Acknowledgments.** This work was partially funded by the Italian Dipartimento della Protezione Civile in the frame of the 2007–2009 Agreement with Istituto Nazionale di Geofisica e Vulcanologia–INGV.

References

Aspinall, W. P., S. C. Loughlin, F. V. Michael, A. D. Miller, G. E. Norton, K. C. Rowley, R. S. J. Sparks, and S. R. Young (2002), The Montserrat Volcano Observatory: Its evolution, organization, role and activities, in *The Eruption of Soufriere Hills Volcano, Montserrat, from 1995 to 1999*, edited by T. H. Druit and B. P. Kokelaar, *Geol. Soc. Mem.*, *21*, 71–91.

Baxter, P. J., W. P. Aspinall, A. Neri, G. Zuccaro, R. J. S. Spence, R. Cioni, and G. Woo (2008), Emergency planning and mitigation at Vesuvius: A new evidence-based approach, *J. Volcanol. Geotherm. Res.*, *178*, 454–473, doi:10.1016/j.jvolgeores.2008.08.015.

Evans, A. W. (2003), Transport fatal accidents and FN-curves: 1967-2001, report for Health and Safety Executive, HSO Books, Suffolk, U. K.

Evans, A. W., and N. Q. Verlander (1997), What is wrong with criterion FN-lines for judging the tolerability of risk?, *Risk Anal.*, *17*, 157–168.

Fischhoff, B., S. Lichtenstein, P. Slovic, S. L. Derby, and R. L. Keeney (1981), *Acceptable Risk*, Cambridge Univ., Press, New York.

Hall, S. S. (2007), Vesuvius countdown, *Natl. Geogr.*, Sept., 114.

Ho, C.-H., E. I. Smith, and D. L. Keenan (2006), Hazard area and probability of volcanic disruption of the proposed high-level radioactive waste repository at Yucca Mountain, Nevada, USA, *Bull. Volcanol.*, *69*, 117–123.

Jaquet, O., R. Carniel, S. Sparks, G. Thompson, R. Namar, and M. Di Cecca (2006), DEVIN: A forecasting approach using stochastic methods applied to the Soufriere Hills Volcano, *J. Volcanol. Geotherm. Res.*, *153*, 97–111.

Kilburn, C. R. J. (2003), Multiscale fracturing as a key to forecasting volcanic eruptions, *J. Volcanol. Geotherm. Res.*, *125*, 271–289.

Martin, A. J., K. Umeda, C. B. Connor, J. N. Weller, D. Zhao, and M. Takahashi (2004), Modeling long-term volcanic hazards through Bayesian inference: An example from the Tohoku volcanic arc, Japan, *J. Geophys. Res.*, *109*, B10208, doi:10.1029/2004JB003201.

Marzocchi, W., and G. Woo (2007), Probabilistic eruption forecasting and the call for an evacuation, *Geophys. Res. Lett.*, *34*, L22310, doi:10.1029/2007GL031922.

Marzocchi, W., and L. Zaccarelli (2006), A quantitative model for the time-size distribution of eruptions, *J. Geophys. Res.*, *111*, B04204, doi:10.1029/2005JB003709.

Marzocchi, W., L. Sandri, P. Gasparini, C. Newhall, and E. Boschi (2004), Quantifying probabilities of volcanic events: The example of volcanic hazard at Mount Vesuvius, *J. Geophys. Res.*, *109*, B11201, doi:10.1029/2004JB003155.

Marzocchi, W., A. Neri, C. G. Newhall, and P. Papale (2007), Probabilistic volcanic hazard and risk assessment, *Eos Trans. AGU*, *88*, 318, doi:10.1029/2007EO320005.

Marzocchi, W., L. Sandri, and J. Selva (2008), BET_EF: A probabilistic tool for long- and short-term eruption forecasting, *Bull. Volcanol.*, *70*, 623–632, doi:10.1007/s00445-007-0157-y.

Mastrolorenzo, G., P. Petrone, L. Pappalardo, and M. Sheridan (2006), The Avellino 3780-yr-B. P. catastrophe as a worst-case scenario for a future eruption at Vesuvius, *Proc. Natl. Acad. Sci. U. S. Am.*, *103*, 4366–4370.

Neri, A., T. Esposti Ongaro, G. Menconi, M. De'Michieli Vitturi, C. Cavazzoni, G. Erbacci, and P. J. Baxter (2007), 4D simulation of explosive eruption dynamics at Vesuvius, *Geophys. Res. Lett.*, *34*, L04309, doi:10.1029/2006GL028597.

Newhall, C. G., and R. P. Hoblitt (2002), Constructing event trees for volcanic crises, *Bull. Volcanol.*, *64*, 3–20, doi:10.1007/s004450100173.

Pareschi, M. T., R. Santacroce, M. Favalli, F. Giannini, M. Bisson, A. Meriggi, and L. Cavarra (2000), Un Gis per il Vesuvio, report, 57 pp., Felici Artigrafiche Ed., Pisa, Italy.

Rosi, M., C. Principe, and R. Vecci (1993), The 1631 Vesuvian eruption: A reconstruction based on historical and stratigraphical data, *J. Volcanol. Geotherm. Res.*, *58*, 151–182.

Sandri, L., W. Marzocchi, and L. Zaccarelli (2004), A new perspective in identifying the precursory patterns of volcanic eruptions, *Bull. Volcanol.*, *66*, 263–275, doi:10.1007/s00445-003-0309-7.

Scandone, R., G. Arganese, and F. Galdi (1993), The evaluation of volcanic risk in the Vesuvian area, *J. Volcanol. Geotherm. Res.*, *58*, 263–271.

Sparks, R. S. J. (2003), Forecasting volcanic eruptions, *Earth Planet. Sci. Lett.*, *210*, 1–15.

Viscusi, W. K. (1992), *Fatal Tradeoffs: Public and Private Responsibilities for Risk*, Oxford Univ. Press, New York.

Voight, B., and R. R. Cornelius (1991), Prospects for eruption prediction in near real-time, *Nature*, *350*, 695–698.

Williams, S. (2001), *Surviving Galeras*, Little, Brown, London.

Woo, G. (2008), Probabilistic criteria for volcano evacuation decision, *Nat. Haz.*, *45*, 87–97.

Woo, G., and W. P. Aspinall (2005), Need for a risk-informed tsunami alert system, *Nature*, *433*, 457.

W. Marzocchi, Istituto Nazionale di Geofisica e Vulcanologia, Roma 1, Via di Vigna Murata 605, I-00143 Roma, Italy. (warner.marzocchi@ingv.it)

G. Woo, Risk Management Solutions, Peninsular House, 30 Monument Street, London EC3R 8NB, UK. (gordon.woo@rms.com)