Ch 2. Quantifying hazard losses

Jonathan Rougier
Department of Mathematics
University of Bristol

Late draft prepared for *Risk and Uncertainty Assessment for Natural Hazards*, J.C. Rougier, R.S.J. Sparks and L.J. Hill eds, 2013, Cambridge, UK: Cambridge University Press.

1. Introduction

This chapter considers the natural hazard Risk Manager as an agent for choosing between interventions that affect the impact of future hazards. As such, it is possible to proceed broadly quantitatively, and to set aside the many less-quantifiable concerns that arise during a period of real-time hazard management. Thus the aim is to describe a framework in which it is possible to address questions such as “Should we build a firebreak at this location, or leave things as they are?”. This chapter only considers one source of uncertainty, which is the inherent uncertainty of the hazard itself, often termed *aleatory* uncertainty. Chapter 3 of this volume considers the extent to which a quantitative analysis can be extended to include more general *epistemic* uncertainties.

The objective is a precise definition of the common notions of natural hazard risk and uncertainty assessment, within the framework of probability. To this end, Section 2 provides a brief justification for the use of probability as the calculus of uncertainty that is appropriate for natural hazards. Section 3 then defines quantitative risk in a limited but precise way. Section 4, the heart of the chapter, considers the three stages in which one passes from the hazard itself to an evaluation of risk, noting that at each stage there are opportunities for the Risk Manager’s intervention. Section 5 considers the implications of using simulation as the primary computational tool for assessing risk, the uncertainties engendered by limits in computing resources, and ways to quantify these uncertainties. Section 6 describes different types of hazard map. Section 7 concludes with a summary.

Note that I have not provided references for natural hazards examples; plenty can be found in the rest of this volume. Most of the references seem to be to statistics books.

2. Why probability?

Hazard losses are uncertain, and therefore a calculus that accommodates uncertainties is necessary to quantify them. The role of quantification within a formally valid framework was discussed Chapter 1 of this volume. This chapter and Chapter 3 advocate the use of probabilities and the probability calculus, and this is in fact the dominant method for quantifying uncertainty, not just in natural hazards, but in almost every endeavour with substantial uncertainty. However, probability is not the only calculus for uncertainty. A standard reference such as Halpern (2003) covers alternative approaches such as lower and upper probabilities, Dempster-Shafer belief functions, possibility measures, ranking functions, and relative likelihood. Therefore this section provides a brief justification for the use of the probability calculus in natural hazards risk assessment.

The first point to note is that no uncertainty calculus can do a complete job of handling uncertainty. Each calculus represents a normative description of how, on the basis of certain axioms and principles, one can make inferences about uncertain events, and adjust those inferences in the light of additional information. The need for such a framework is attested by the fact that people are demonstrably not good at handling uncertainty in even quite simple situations. Gigerenzer (2003) provides an enjoyable and informative introduction, while Lindley (1985) provides an introduction to the use of probability in reasoning and decision-making.

To illustrate, consider the axioms and principles of the probability calculus. The three axioms concern probabilities defined on events (or propositions) that are either true or false. They are:

1. All probabilities are non-negative;
2. The probability of the certain event is 1;
3. If events A and B cannot both be true, then Pr(A or B) = Pr(A) + Pr(B).

These axioms, and their generalisations, are discussed in most textbooks in probability and statistics; see, e.g., Grimmett and Stirzaker (2001), chapter 1, or DeGroot and Schervish (2002), chapter 1.

There are really two principles in probability. The first is a principle for extending probability assessments from a limited set of events to a more complete set that is appropriate for the information that might be collected, and the decisions that must be made. In fact, formal treatments using the probability calculus usually sidestep this issue, by starting from the notion that probabilities have already been assigned to all events in a *field*. A field is a collection of sets of events, including the null event, which is closed under complements and unions. The assertion is that this field is appropriate in the sense given above; i.e. it contains all the events of interest. But there are more general approaches to extending probability assessments, discussed, for example, in Paris (1994, chapter 6).

The second principle of the probability calculus is that inferences are adjusted by conditioning, so that if event B were found to be true, then the adjusted probability of A would be Pr(A and B) / Pr(B). The axioms and the principle of conditionalisation have several different justifications, including: probabilities as relative frequencies, as subjective degrees of belief, as logical consequences of more primitive axioms for reasoning. These are outlined in Hacking (2001), with a much more forensic but technical assessment in Walley (1991), especially chapters 1 and 5.

Halpern (2003, page 24) lists the three most serious problems with probability as (i) probability is not good at representing ignorance, (ii) an agent may not be willing to assign probabilities to all events, and (iii) probability calculations can be expensive. In the context of natural hazards, I think (ii) is the most concerning, and (iii) can be an issue, although is becoming less so as statistical computational techniques continue to improve, along with computer power. (i), as Halpern himself notes, can be handled (in principle and at least partially in practice as well) within the probability calculus by extending the set of events. This is effectively the approach explored in Chapter 3 of this volume.

All uncertainty calculuses have axioms and principles. In each case, the agent will find herself asking, of the axioms and principles: Do I believe that? Should I believe that? These questions do not have simple answers, which should not be surprising in the light of the proliferation of competing approaches to uncertainty representation. Consequently, an agent does not sign up for a particular uncertainty calculus with the conviction that it is exactly what is required for her situation. It is very important for her to understand the limitations of her chosen calculus, so as not to overstate the result, and to have methods for addressing the limitations informally.

In natural hazards things become more complicated, because the agent is the Risk Manager, who must satisfy an Auditor. What seems to be a compelling calculus to the Risk Manager, possibly after detailed study and reflection, may seem much less so to the Auditor, and to the stakeholders he represents. Overall, this line of reasoning suggests favouring calculuses with simple axioms, and easily-stated limitations.

On this basis, the probability calculus is highly favoured: its axioms are the simplest, and its main limitation is that it is necessary to specify probabilities for all events in an appropriate field. In practice this limitation can be addressed informally using sensitivity analysis, in which alternative choices for hard-to-specify probability distributions are tried out. The tractability of probability calculations makes sensitivity analysis a feasible strategy for all but the largest problems. And the universality of probabilistic reasoning makes communication relatively straightforward. That is not to say, though, that these issues are clear-cut. It is important for all parties to appreciate that not all aspects of uncertainty can be expressed as probabilities, and that a probabilistic analysis may create the false impression that all uncertainty has been accounted for.

3. A quantitative definition of risk

Risk is a multivalent and multivariate concept, yet it is often treated as though it was a well-defined scalar, for example in statements comparing the 'riskiness' of different hazards. This indicates the nature of quantified risk as a summary measure whose role is to represent the gross features of a hazard, and to operate at the first stage of a 'triage' of hazards and actions.[1] In their review of risk terminology, the UK Central Science Laboratory (CSL) recommends that the definition of risk "should include both probability and the degree of effect, including its severity, but in a way that keeps them distinct and gives rise to a single dimension." (Hardy *et al.*, 2007, page 70). The presumption in this recommendation is that uncertainty be quantified in terms of probabilities, as discussed in the previous section. I will adopt 'loss' as the portmanteau term for the quantifiable aspects of harm and damage that follow from a natural hazard event.

[1] 'Triage' in the sense of sorting by priority and expediency, as in the Emergency Room of a hospital.

To illustrate, supervolcanoes are more risky than asteroid strikes because the loss in the two cases is the same order of magnitude (i.e. catastrophic), but the probability of a supervolcano is higher than an asteroid strike on human time-scales (Mason *et al.*, 2004).

If risk is a summary term then it should be derived from a more detailed analysis. The second stage of a triage is then to refer back to this detailed analysis for those hazards for which such an assessment is required. To my mind, this constrains the operational definition of risk to be the mathematical expectation of loss, and the detailed analysis for which it is the summary is the loss probability distribution function. The usual feature of this distribution for natural hazards would be a long right-hand tail (positive skewness), indicating that small losses are common, but occasionally very large losses occur. If X denotes the unknown loss that a hazard induces over a specified time interval, then its distribution function is denoted

$$F_X(x) = \Pr(X \leq x),$$

where 'Pr' denotes probability, and 'x' denotes the abscissa. This distribution function is usually visualised in terms of the loss Exceedance Probability (EP) curve, which is the graph of $1 - F_X(x)$ on the vertical axis against the ordinate x on the horizontal axis, as shown in Figure 2 (in statistics this would be called the 'survivor function'). This graph does not necessarily start at (0, 1); in fact it starts at (0, 1-p) where p is the probability that there are no losses during the time interval. The EP curve must attain zero at some finite value for x because there is always a limit to how much can be lost; therefore, having bounded support, X must have finite moments, and so its mathematical expectation and variance certainly exist.

[FIGURE 1 ABOUT HERE]

Risk may then be defined as the mathematical expectation of loss, i.e. the sum of the product of each possible loss amount and its probability. This is also referred to as the 'mean loss'. This definition has the features outlined in the CSL recommendation above: it is a scalar quantity measured in the same units as the loss, and incorporates both the loss estimates and their probabilities. The same sized risk could indicate a hazard that will probably induce a medium-sized loss, or a hazard that will probably induce a small loss but occasionally induce a very large loss. These two cases cannot be distinguished in a scalar summary, but they can be distinguished in the underlying probability distributions.

Within statistics this definition of risk as expected loss has a long provenance in the field of Decision Theory; see, e.g., Rice (1995), Chapter 15. In Catastrophe Modelling for insurance, the EP curve and risk focus on financial loss for a one-year time interval. The risk is termed the Average Annual Loss (AAL), and represents the 'fair price' for an insurance premium.

There is a very strong connection between this definition of risk and the EP curve: *risk is the area under the EP curve*. The fact that the expected loss is equal to the area under the EP curve is a purely mathematical result (see, e.g., Grimmett and Stirzaker, 2001, page 93), but it is, from our point of view, a very useful one as well. It indicates that the gross comparison of different hazards, or of different actions for the same hazard, can be done by plotting their EP curves on the same graph and comparing the area underneath them. Then, in cases where this gross comparison is not sufficient, a more detailed assessment of the EP curve can be made, comparing, for example, the probabilities of very large losses.

Summary. The loss due to a hazard for a specified time interval is an uncertain quantity that is represented by a probability distribution; this is usually visualised in terms of an Exceedance Probability (EP) curve. Risk is defined to be the expected (or mean) loss; it is a mathematical result that this is equal to the area under the EP curve.

4. Structural modelling of hazard outcomes

This section follows the flow of information and judgement from the hazard to the loss. It proceeds in three stages: representing the aleatory uncertainty of hazard outcomes, representing the ‘footprint’ of a hazard outcome in space and time, and representing the loss of this hazard outcome as a loss operator applied to the footprint. Probabilities assigned to hazard outcomes ‘cascade’ through the stages to induce a probability distribution for loss.

It is worth pausing to examine the need for such a tripartite representation. After all, it would be much easier simply to construct a catalogue of ‘similar’ hazard events, and then use the collection of losses from those events as a proxy for the loss distribution. Why bother with all the extra modelling?

Two of the reasons are driven by the need for the scientific analysis to inform the Risk Manager’s decisions: non-stationary in the hazard domain, and the intention to intervene. The natural hazard Risk Manager has to operate on time intervals of thirty years or more, over which many aspects of the hazard loss can be expected to change substantially; this is the problem of *non-stationarity*. For example, increasing populations result in increasing city size, increasing population density in cities, changes in the quality of buildings and infrastructure, changes in land-surface characteristics in catchments. In this situation historical losses are not a reliable guide to future losses over the whole of the time interval, and additional judgements are needed to perform the extrapolation.

Second, the Risk Manager is concerned explicitly with evaluating actions, which represent *interventions* to change the loss that follows a hazard event. For example: constructing defenses, changing building regulations, or relocating people living in hazardous areas. For this purpose, it is important to model the hazard in a causal rather than a statistical way. The key role of the tripartite representation is that it provides a framework within which interventions can be modelled and compared. Each possible intervention generates its own EP curve and its own risk. Then the Risk Manager has the task of choosing between interventions according to their EP curves, taking account of the costs and benefits of intervening and the probability of completing the intervention successfully.

A third reason is that the catalogue of observed hazard events is likely to be incomplete. This is a concern shared by insurers, who otherwise operate on short time intervals and are not concerned with interventions. There is a large probability that the next event will be unlike the catalogued ones, and so physical insights are required to extrapolate from the catalogue in order to build up a picture of what might happen. In some ways this is also an issue of non-stationarity. The catalogue can be extended by widening the criteria under which events can be included. In earthquake catalogues, for example, the events observed at a specified location may be augmented with events observed at other similar locations. Effectively this makes a judgement of spatial and temporal stationary; hence non-stationarity is one reason why catalogues might be very incomplete.

It is questionable to what extent any framework can encompass all features of natural hazards. The framework presented here seems sufficiently general, even if some 'shoe-horning' is necessary for particular hazards. It is sufficient to illustrate the main concepts, and there is also value in considering how well a particular hazard does or does not fit. It also corresponds closely with the modular approach adopted by catastrophe modelling companies.

4.1. The hazard process

The specification of the hazard process comprises three stages:

1. Specification of the hazard domain;
2. Enumeration of the hazard events; and
3. Assignment of probabilities to hazard outcomes.

In Step 1, the hazard domain is simply the spatial region and the time interval over which the hazard is being considered. This has to be specified *a priori*, in order that the probabilities in Step 3 are appropriately scaled, particularly over the time interval. It would be common for the time interval to start 'now'. As already noted, the Risk Manager is likely to be considering time intervals of many years: thirty, for example, for a hazard such as an earthquake, longer for nuclear power installations, and longer again for nuclear waste repositories.

In Step 2, each hazard event is described in terms of a tuple (a tuple is an ordered collection of values). This tuple will always include the inception time of the event, therefore it is written (t, ω), where ω is a tuple that describes the hazard event that starts at time t. For an earthquake, ω might comprise an epicentre, a focal depth, and a magnitude; or, perhaps, a marked point process representing a local sequence of shocks with specified epicentres, focal depths, and magnitudes. For a rain storm, ω might comprise a time-series of precipitations, where each component in the time-series might itself be a spatial map. Likewise with a hurricane, although in this case each component of the time-series might be a spatial map of wind velocities. The set of all possible values of ω is denoted as Ω.

In Step 3 it is important to distinguish between a *hazard event,* and a *hazard outcome*. The hazard outcome is the collection of hazard events that occur in the hazard domain. Therefore {(t, ω), (t', ω')} would be a hazard outcome, comprising two events: at time t, ω happened, and at time t', ω' happened. The collection of all possible hazard outcomes is very large, and, from the Risk Manager's point of view, each outcome must be assigned a probability. In practice this assignment of probabilities can be 'tamed' as follows, recollecting that ω describes all features of the hazard event bar its inception time (i.e. including location and magnitude):

Simplifying choices:

1. Different ω correspond to probabilistically independent processes, and
2. Each process is a homogeneous Poisson process with specified rate λ_ω.

Under these simplifying choices, the probability of any hazard outcome can be computed explicitly, once the rates have been specified. The simplifying choices imply that the time between hazard events is Exponentially distributed with a rate equal to $\lambda = \sum_\omega \lambda_\omega$, and that the probability that the next event is ω is equal to λ_ω / λ; see, e.g., Davison (2003), examples 2.35 and 2.36.

Interventions. The Risk Manager has few opportunities to intervene in the hazard process, which in this case would take the form of changing the probabilities on the hazard outcomes. Examples would be controlled burning for forest fires, and controlled avalanches. In situations where the trigger is sometimes man-made, the probabilities can be reduced by information campaigns, regulations, and physical exclusion. Again, this applies mainly to forest fires and avalanches; also perhaps to terrorism. One can also consider more speculative possibilities; for example, geo-engineering solutions such as cloud-seeding, to reduce the probability of large rain storms over land.

4.2. The footprint function

The footprint function represents the 'imprint' of a hazard outcome at all locations and times in the hazard domain. The purpose of the distinction between the hazard outcome and its footprint is to separate the hazard outcome, which is treated as *a priori* uncertain, from its effect, which is determined *a posteriori* largely by physical considerations. For example, for inland flooding the hazard outcome would describe a sequence of storm events over a catchment. The footprint function would turn that sequence of storm events into a sequence of time-evolving maps of water flows and water levels in the catchment.

Often the footprint function will be expressed in terms of an initial condition and a hazard event. The initial condition describes the state of the hazard's spatial domain just prior to the hazard event. In the case of inland flooding, this might include the saturation of the catchment, the level of the reservoirs and rivers, and possibly the settings of adjustable flood defenses such as sluice gates. The footprint function of a hazard *outcome* is then the concatenation of all of the hazard event footprint functions, where the initial condition for the first event depends on its inception time, and the initial conditions for the second and subsequent events depend on their inception times, and the hazard events that have gone before. This is more sophisticated than current practice, which tends to be event-focused, but it shows that there is no difficulty, in principle, to generalising to hazard outcomes that comprise multiple events; nor, indeed, to outcomes that involve multiple hazards.

Most natural hazards modelling currently uses deterministic footprint functions. For flows (volcanoes, landslides, avalanches, tsunamis and coastal flooding) these are often based on the shallow water equations, single or multi-phase, or, for homogeneous flows, sliding block models. Earthquakes use the equations of wave propagation through elastic media. For inland flooding, there is a range of different footprint functions, from topographically explicit models, through compartmental models, to empirical models. Sometimes physical insights are used to derive parametric relationships between a hazard event and particular features of the hazard footprint, and the parameters are then statistically fitted to a catalogue of similar events. Woo

(1999) contains many examples of of these types of relationships, based on dimensional or scaling arguments (often implying power laws).

In this chapter we treat the footprint function (and the loss operators, below) as known. In practice they are imperfectly known, which is an important source of epistemic uncertainty, discussed in Chapter 3 of this volume.

Interventions. The footprint function provides an important route through which the Risk Manager can model the effect of interventions. These would be interventions that interfered physically with the impact of the hazard on the region. For example, building a fire-break, a snow dam, or a levée. Typically, these interventions change the topography and other land-surface features local to the hazard event, and thus change the evolution of the event's impact in space and time; for example, diverting a flow, or retarding it, to allow more time for evacuation.

4.3. The loss operator

The footprint of the hazard is neutral with respect to losses. Loss is a subjective quantity, and different Risk Managers with different constituencies will have different notions of loss, although these will typically centre on loss of life or limb, loss of ecosystem services, and financial loss due to damage to property. The purpose of the functional separation between the footprint function and the loss operator is to allow for this distinction between what is generic (the hazard footprint) and what is subjective (the Risk Manager's loss). Any particular loss operator transforms a hazard outcome's footprint into a scalar quantity. Typically, this will take the form of the addition of the losses for each hazard event, where it is likely that the losses from later events will depend on the earlier ones. For simplicity I will treat the loss operator as deterministic, but the generalisation to an uncertain loss operator is straightforward, and discussed in Chapter 3 of this volume.

The nature of the loss operator depends on the hazard and the Risk Manager. Consider, as an illustration, that the hazard is an earthquake and the Risk Manager represents an insurance company. The buildings in the insurance portfolio are each identified by location, type, and value. The type indicates the amount of ground acceleration the building can withstand, and a simple rule might be that accelerations above this amount destroy the building, while accelerations at or below this amount leave the building intact. To compute the loss of an event, the footprint function is summarised in terms of a spatial map of peak ground acceleration, and this would be converted into a loss (in millions of dollars) by determining from this map which of the insured properties were destroyed. (Naturally, this is a rather simplified account of what actually happens.)

In the same situation, if the Risk Manager's loss operator concerns loss of life, then the same peak ground acceleration map from the footprint function might be combined with maps of building type and population density. A peak ground acceleration map is an example of an *extremal hazard map*, discussed in Section 6.

Interventions. Interventions in the loss operator concern changes in vulnerability. For example, the Risk Manager for the insurance company can change the loss distribution by modifying the premiums for different types of property, or by refusing to insure certain types of building. The Risk Manager concerned with loss of life could change building regulations, or rezone the city to move people away from areas with very high peak ground accelerations. Another example of an intervention to affect the loss operator is the installation of an early warning system, e.g. for tsunamis, or the implementation of a phased alert scheme, e.g. for wildfires. This might be accompanied by a public education programme.

4.4. Summary

The chain from hazard to the distribution of loss has three components. First, the hazard process itself, which comprises the enumeration of different hazard events, and then the assignment of probabilities to hazard outcomes. The distinction between hazard events and hazard outcomes is crucial (see Section 4.1). Second, the footprint function, which represents a hazard outcome in terms of its impact on the hazard domain. Third, the loss operator, which varies between Risk Managers, and maps the hazard footprint into a scalar measure of loss. The Risk Manager has the opportunity to intervene in all three components: by changing the probabilities of the hazard outcomes, by changing the local topography and therefore the hazard footprint, or by changing the vulnerability, and therefore the loss operator. This is summarised in Table 1.

Table 1. Summary of the three components of the hazard risk assessment, with key concepts and opportunities for the Risk Manager's intervention.

	Concepts	**Opportunities for intervention**
Hazard process	Hazard's spatial-temporal domain, events and outcomes, probabilities, simplifying choices	Reducing probability of human triggers: wildfires, avalanches
Footprint function	Objective, typically implementing physical equations or statistical regularities. Shows the imprint of the hazard outcome in space and time, often in terms of individual hazard events.	Change the topography of hazard's spatial domain: levées, dams, firebreaks
Loss operators	Subjective, depending on constituency and Risk Manager. Typically measuring loss of life, loss of ecosystem services, or financial damage	Reduce vulnerability: regulations, rezoning and relocation, harden critical infrastructure.

5. Estimating the EP curve

The EP curve (see Section 3) represents the distribution function of the Risk Manager's loss. There is a standard formula for computing the EP curve, which is simply to sum the probabilities for all outcomes which give rise to a loss greater than x, for each value of x. In practice, this calculation would usually be done by simulation. This has important implications for how the EP curve and quantities derived from it are reported, particularly high percentiles.

5.1. Uncertainty and variability

There is a very important distinction in statistics between *uncertainty* and *variability* (see, e.g., Cox, 2006, notably chapter 5), and a failure to understand this distinction lies at the heart of many suspect attempts to quantify uncertainty using variability. Uncertainty about the loss induced by a hazard is represented in the form of a probability distribution function, the function denoted F_X in Section 3. From a probabilistic point of view, this probability distribution function is a complete description of uncertainty. In practice, we are aware that when we construct such a function we are obliged to introduce various simplifications. In this Chapter we have considered the aleatory uncertainty of the hazard outcomes as the *only* source of uncertainty. This is very great simplification! In Chapter 3 of this volume, we extend our consideration to incorporate epistemic sources of uncertainty, such as incomplete information about the hazard outcome probabilities, or about the footprint function, or about the loss operator(s). The effect of this extension is to make a better assessment of the probability distribution function of loss. Or, as discussed in Chapter 3, to produce several different probability distributions, each one specified conditionally on certain simplifications.

But representing our uncertainties probabilistically is not the only challenge. The computational framework, in which uncertain hazard outcomes cascade through the footprint function and the loss operator(s), can always be simulated, but often this simulation will be expensive. With limited resources, we will have an incomplete knowledge of the probability distribution of loss based upon the simulations we have, say n simulations in total. To put this another way, were we to have done the n simulations again with a different random seed, the sample of losses would have been different, and the summary statistics, such as the expected loss, would have different values. This is the problem of *variability*. Expensive simulations mean that our knowledge of the loss distribution is limited, and that our numerical descriptions of it are only approximate.

In statistics, variability is summarised by *confidence intervals*. It is very important to appreciate that a confidence interval is *not* a repesentation of uncertainty. The distribution function F_X is a representation of uncertainty. A confidence interval is a representation of variability of our estimate of some feature of F_X (like its mean), which reflects the fact that n, the number of simulations we can afford to do in order to learn about this feature, is not infinite. To state that the risk (i.e. mathematical expectation) of a hazard is $20.9M is a statement about uncertainty. To state that the 95% confidence interval for the risk of a hazard is [$17.1M, $25.3M] is a statement about both uncertainty and variability: the location of this interval is an assessment of uncertainty, while its width is an assessment of variability. The latter will go down as n increases.

I will focus here on the variability that arises from not being able to perform a large number of natural hazards simulations. But in some situations, variability can also arise from the limited information about the hazard itself. For example, a flood engineer might use the historical distribution of annual maximum river heights at some specified location to estimate the probabilities in the hazard process. If the record does not go back very far, or if only the recent past is thought to be relevant because of changes in the environment, then estimates of the probabilities based on treating the historical measurements as probabilistically independent realisations from the hazard process will be imprecise. I prefer to treat this as a manifestation of epistemic uncertainty about the hazard process, and this is discussed in Chapter 3 of this volume.

This Section is about assessing confidence intervals for properties of the loss distribution. My contention is that variability ought to be assessed because natural hazards loss simulations are expensive, and some important assessments of uncertainty, such as high percentiles of the loss distribution, will tend to be highly variable.

5.2. Monte Carlo simulation and EP curve variability

By simulating hazard outcomes we can estimate the EP curve; technically, simulation is an implementation of Monte Carlo integration (see, e.g., Evans and Swartz, 2000). Better estimates, i.e. ones that tend to be closer to the true value, can be achieved at the expense of a larger number of simulations, or a more careful experimental design (e.g., applying variance reduction techniques such as control variables, or antithetic variables). For complex applications, where resources constrain the number of simulations, measures of variability are required to indicate the closeness of the estimated EP curve to the true EP curve, and likewise for other quantities related to the EP curve, like the mean, or the 95th or 99.5th percentiles of the loss distribution.

A Monte Carlo simulation is constructed for a specific hazard process, footprint function, and loss operator: uncertainty about these is the subject of Chapter 3 in this volume. It has the following steps:

1. Simulate a hazard outcome,
2. Evaluate the footprint function for the outcome,
3. Evaluate the loss operator for the footprint.

Repeating this many times for probabilistically independent simulations of the hazard outcome will build up histogram of losses, which is an estimator of the probability distribution of losses. If the footprint function or loss operator are stochastic, then they too can be simulated rather than simply evaluated.

Step 1 can be complicated, bearing in mind the distinction between a hazard event and a hazard outcome: the latter being the concatenation of many events over the hazard domain. But accepting the simplifying choices described in Section 4.1 makes this step straightforward. The sequence of events that make up a hazard outcome may also affect the footprint function and the loss operator. Simplifying choices in the same vein would be to treat the interval between hazard events as sufficiently long for the hazard domain to ‘reset’, so that the footprint and loss of later events does not depend on earlier ones. This seems to be a common choice in catastrophe modelling.

The variability of a Monte Carlo estimator can be quantified in terms of confidence intervals or confidence bands. A 95% confidence band for the EP curve comprises a lower and upper curve with the property that the true EP curve will lie entirely within these curves at least 95% of the time. Confidence bands for the EP curve can be computed using the Dvoretzky-Kiefer-Wolfowitz Inequality, as described in Wasserman (2004), chapter 7. For a 1-α confidence band the upper and lower curves are ± the square root of ln(2/α) / (2n) vertically about the empirical EP curve, where n is the number of simulations. For example, with n = 1000, a 95% confidence band is ±0.043. Figure 2 shows estimated EP curves and 95% confidence bands for different numbers of simulations.

[FIGURE 2 ABOUT HERE]

5.3. Estimating the risk

The risk is defined to be the expected loss, as previously discussed. This is easily estimated from a random sample, and a Normal confidence interval will likely suffice, or a Student-t interval if the number of simulations is small. For the four simulations given above, the estimated risk and a 95% CI is: n = 10, $9.8M (5.2, 14.4); n = 100 $9.2M (7.6, 10.8); n = 1000, $9.7M (9.2, 10.2); n = 10000, $9.5M (9.4, 9.7). With 10000 simulations, the proportional uncertainty is only a few percent.

5.4. Estimating high percentiles

The Solvency II Directive[2] on Solvency Capital Requirement is described in terms of the 99.5th percentile of the loss distribution for a one-year time interval (Paragraph 64, page 7). In the section 'Statistical quality standards' (Article 121, page 58) there is no guidance about how accurately this value needs to be estimated, and the issue that the 95% confidence interval for the 99.5th percentile may be very large is not explicitly addressed, although it has certainly been discussed by catastrophe modelling companies. Here I outline briefly how such an interval might be assessed.

First, an obvious point: estimators of the 99.5th percentile based on less than at least n = 1000 simulations will tend to be unreliable. For smaller n it will be very hard to quantify the variability of these estimators. The technical reason is that convergence of the sampling distribution for estimators of high percentiles to a form which admits a pivotal statistic is slow. Therefore confidence intervals based on asymptotic properties (as most confidence intervals will be, in practice) may have actual coverage quite different from their notional coverage of 1-α. Thus, while one can certainly estimate the 99.5th percentile with a small simulation, and indeed compute a confidence interval, an Auditor should treat these with caution.

So how might we proceed in the case where n is at least 1000, and preferably much larger? One possibility for very high percentiles is to use a parametric approach based on Extreme

[2] http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2009:335:0001:0155:EN:PDF

Value Theory; see, e.g., de Haan and Ferreira (2006), Chapter 4—this is highly technical. Often a more transparent approach will suffice. A simple estimator of Q, the desired percentile, is to interpolate the ordered values of the simulated losses. A 95% confidence interval for Q based on this estimator can be assessed using the Bootstrap. By a loose analogy with the Exponential distribution, a nonparametric Bootstrap on the logarithm of Q should be effective; see Davison and Hinkley (1997), Sections 2.7 (notably Example 2.19) and 5.2. But one must bear in mind that the Bootstrap is a very general technique with lots of opportunities for both good and bad judgements. It would be wise to treat the bootstrap 95% confidence interval of the 99.5th percentile as merely indicative, unless the number of simulations is huge (tens of thousands).

For the illustration, the estimated 99.5th percentile and approximate 95% confidence intervals are: n = 1000, \$43.8M (36.9, 53.8); n = 5000, \$42.3M (40.5, 45.2), n = 10000, \$41.7M (40.3, 43.4). Here, even 10000 simulations gives a confidence interval with a width of several million dollars, which is a proportional uncertainty of about ten percent.

6. A digression on terminology

In natural hazards, one often finds statements such as "ω is an event with a return period of k years", where k might be 100 or 200, or 1000. Or, similarly, "ω is a one in k year event". Personally, I do not find such terminology helpful, because it is an incomplete description of the event, and, in order to be useful, must apply within a strongly parameterised hazard process. I would rather see these issues made explicit.

The explicit statement would be "Considering the hazard process over the next w years, the marginal distribution of event ω is a homogeneous Poisson process with arrival rate 1/k." This makes the time interval over which the description holds explicit. It also translates the absence of any further information into an explicit statement of homogeneity. If such homogeneity were not appropriate, then more information should have been given in the first place. The fact that such a statement is not qualified by any reference to other events suggests that the conditional distribution of ω is judged to be unaffected by the other events. Hence such statements embody completely the simplifying choices of section 4.1.

Perhaps we can take it as conventional that all statements about return periods embody these simplifying choices. But what is still missing in the original statement is an explicit time domain: for how long into the future does event ω have an arrival rate of 1/k? There is no convention about such an interval, and so all such statements need to be qualified. But then we are in the realm of "ω is an event with a return period of k years, over the next w years". Many people will find two references to different numbers of years confusing. But considering that this is a more precise statement than before, one wonders how much confusion there is currently. Do people, for example householders, perhaps think that a return period of 1000 years is good for the next 1000 years? Or only for the next year? Or for the next few years? This is a source of needless ambiguity that can be totally eliminated with a little more precision.

The same language crops up in another context as well. The “one in k year loss” is synonymous with the 1 - 1/k^{th} quantile of the one year loss distribution. So why not say this? Rather than the “one in 200 year loss”, state “99.5th percentile of the one year loss distribution”.

Personally, I think this evasive language reflects a reluctance to use explicit probabilistic constructions. Such constructions would show clearly how subjective these assessments are, depending as they do on very strong judgements that would be hard to test with the limited data available. The simplifying choices for the hazard process are very strong; and one does not get to the 99.5th percentile of the loss distribution without also making some strong judgements about distributional shapes. But concealing this subjectivity is the wrong response, if our intention is to inform the Risk Manager. Rather, we should get it out in the open, where it can be discussed and refined. In some cases we will be forced into strong judgements for reasons of tractability. In this case, too, it is necessary to acknowledge these judgements explicitly, so that we can decide how much of a margin for error to include in the final assessment. This is discussed in more detail in Chapter 3 of this volume.

7. Creating hazard maps

There is no general definition of a hazard map. Although the term tends to have recognised meaning within individual natural hazard fields, methods for producing such maps vary widely. One important distinction is whether the map shows features of the hazard footprint, or of the hazard loss. For losses, it makes sense to add losses over events (although the size of the loss from later events may depend on the earlier ones). This is not true for aspects of the hazard footprint, even though these may relate strongly with losses per event.

For example, for an earthquake event, peak ground acceleration at a given location is often taken to be strongly related to damage to buildings at that location. But the addition of peak ground accelerations at a given location across the events in a hazard outcome does not largely determine the loss from a hazard outcome at that location: many small shocks do not incur the same loss as one large one, even if later losses are not affected by earlier ones.

Therefore I favour reserving *risk map* for showing losses, in situations where the loss over the hazard domain is the sum of the losses over each location. In this case, a risk map would show the expected loss at each location, and the total risk of the hazard would be the sum of the risks in each location. (Technically, this follows from the linear property of expectations: if risk were not defined as an expectation, there would be no reason for risk maps to have this appealing property.) In other words, *risk maps summarise the losses from hazard outcomes*.

What about maps that summarise probabilistic aspects of the hazard footprint, for which *(probabilistic) hazard map* seems appropriate? Even within a single hazard, such as volcanoes, there are a number of different ways of visualising the uncertain spatial footprint of the hazard. This variety seems strange, because in fact it is very clear what features a probabilistic hazard map ought to have. These are summarised in Table 2.

Table 2: Features to be found in a probabilistic hazard map.

1. Clearly stated time interval over which the probabilities apply.
2. Clearly stated extremal operator, which summarises hazard outcomes over the time index.
3. Clearly stated threshold for which probabilities of exceedance will be computed in each pixel.
4. Contours or shading indicating regions of similar probability.

The first stage of constructing a hazard map is to reduce the hazard footprint to a spatial map. This means summarising over the time index, if one is present. A standard summary would be to take the maximum of some component over time, and for this reason I refer to the resulting maps as *extremal hazard maps:* note that these are constructed for a specific outcome (or event). Whatever summary is taken, the general principle for informing the Risk Manager is to use the value that relates best to loss at the location, and this will typically be a maximum value. For flooding, for example, the summary might be the maximum depth of inundation over the hazard outcome, at each location. An illustration for an earthquake is shown in Panel A of Figure 3.

[FIGURE 3 ABOUT HERE]

The second stage is then to attach probabilistic information from the hazard process to each extremal hazard map. If the extremal hazard maps show hazard outcomes then attaching probabilities is straightforward, since we have assigned probabilities directly to hazard outcomes over the hazard domain. Likewise, if the simplifying choices have been used, then any particular concatenation of events into a hazard outcome can be assigned an explicit probability based on the rates of each outcome. Other situations, in which more complicated hazard processes have been used, or where the footprint of a later event depends on the earlier ones, will probably have to be handed through simulation.

The third stage is to display the resulting set of probability-weighted extremal hazard maps as a single map. Each pixel has its own distribution for the summary value, but this is too much information to show on one map (although a visualisation tool would give the user the option to click on a pixel and see a distribution). Therefore further information reduction is required. The most useful approach, in terms of staying close to losses, seems to be to construct a map showing the probability of exceedance of some specified threshold, at each location; this is what I term a *probabilistic hazard map*. It will often be possible to identify a threshold at which losses become serious, and then this map can be interpreted loosely as ‘probability of serious loss from the hazard over the next w years’, where w is specified. Such a map identifies *hazard* *zones*, but it would be incorrect, in my treatment, to call them risk zones (because there is no explicit representation of loss). An illustration of a probabilistic hazard map for earthquakes is shown in Panel B of Figure 3.

UK Environment Agency flood maps. What would one expect to find in a probabilistic flood hazard map? I would expect the map to reflect a specified time interval that was relevant to Risk managers and households, something between one to five years (or else produce maps for different time intervals), and to show the probability of inundation at each location; i.e. a threshold set at zero inches, or maybe slightly more to allow for some imprecision. Then these probabilities could be displayed using a simple colour scale of, say, white for probability less than 10^{-3}, light blue for probability less than 10^{-2}, dark blue for probability less than 10^{-1}, and red or pink for probability not less than 10^{-1}.

The UK Environment Agency produces flood maps that are superficially similar to this, at least in appearance.[3] The threshold seems to be zero inches, and there are dark blue areas around rivers and coasts, light blue areas outside these, and white elsewhere (there is no pink—perhaps that would be too scary). There is a distinction between the probabilities of coastal and inland flooding, but this is largely immaterial given the scale of the colour scheme. However, the checklist in Table 2 reveals some concerns. No time interval is specified, and in fact the zones do not show probabilities, but instead the extent of 1 in 100 year (probability = 0.01 in one year, dark blue) and 1 in 1000 year (probability = 0.001 in one year, light blue) events. This focus on events rather than outcomes makes it impossible to infer probabilities for hazard outcomes for a specific time interval except under the very restrictive condition *that these are the only two flooding events that can occur*. In this case, the interpretation in section 6 would allow us to convert the EA flood map into a probabilistic hazard map for any specified time interval. But of course this is a totally indefensible condition.

The origin and development of the EA flood maps is complicated; see, for example, Porter (2009). They were never meant to be probabilistic hazard maps, and this focus on two particular events was driven in part by the requirements of Planning Policy Statement 25.[4] Unfortunately, from the point of view of risk management, town planners, actuaries, businesses, and householders cannot inspect the map and infer a probability for flood inundation over a specified period. But a probabilistic hazard map would have been much more expensive to compute, because a much wider range of events would have had to have been assimilated.

8. Summary

Ambiguity and imprecision are unavoidable when considering complex systems, such as natural hazards and their impacts. That is not to say, though, that one cannot be systematic in developing a formal treatment that would serve the needs of the Risk Manager. Such a treatment removes needless ambiguity by the use of a controlled vocabulary. Natural hazards, for all their diversity, show enough common features to warrant a common controlled

[3] http://www.environment-agency.gov.uk/homeandleisure/37837.aspx

[4] http://www.communities.gov.uk/documents/planningandbuilding/pdf/planningpolicystatement25.pdf

vocabulary. This Chapter is an attempt to specify such a vocabulary within a probabilistic framework, defining: hazard domain, hazard event, hazard outcome, footprint function, loss operator, EP curve, risk, extremal hazard map, probabilistic hazard map, and risk map. The entry points where judgements are required are the probabilities of the hazard outcomes, the footprint function, and the loss operator(s). All other aspects of the hazards analysis presented here then follow automatically, and their production is effectively an issue of statistical technique and computation.

Two issues are worth highlighting. First, there is a crucial distinction between hazard events and hazard outcomes, which encompasses both the distinction between one event and many, and between a single hazard type and different hazard types. This follows from the specification of a hazard domain in which several different hazard events spanning more than one hazard type might occur. This is central to the role of the Risk Manager, who is concerned with losses *per se*, but less so to the insurer, who is able through a contract to limit losses to certain hazard types. This distinction infiltrates other aspects of the analysis; for example the EP curve is a description of outcomes, not events, likewise risk maps and hazard maps as defined here. The probabilistic link between events and outcomes is complicated, but can be tamed using the two simplifying choices described in Section 4.1. Similarly, footprint functions and loss operators are defined on outcomes not events, but there are obvious simplifications that allow definitions on events to be extended to outcomes.

Second, the use of simulation to construct EP curves and their related quantities should be seen as an exercise in statistical estimation. Hence EP curves should be presented with confidence bands to assess variability, and risk and quantiles should be presented with confidence intervals. The high percentiles required by regulators, e.g. 99.5th percentile for Solvency II, present serious problems for estimation, and the confidence intervals of such percentiles may be both wide and unreliable.

It is also important to reiterate the point made in Section 2. The probability calculus seems to be a good choice for quantifying uncertainty and risk in natural hazards. But it is by no means perfect, and nor is it the only choice. Sensitivity analysis with respect to hard-to-specify probabilities is crucial when assessing the robustness of the analysis. The importance of sensitivity analysis favours a simpler framework that can easily be replicated across different choices for the probabilities (and also for other aspects such as structural and parametric choices in the footprint function, see Chapter 3 of this volume), over a more complicated framework which is too expensive to be run more than a handful of times. From the point of view of the Risk Manager, and bearing in mind how much uncertainty is involved in natural hazards, I would favour simpler frameworks that function as tools, rather than a more complicated framework that we are forced to accept as 'truth'. This issue pervades environmental science; Salt (2008) is a good reality check for modellers who may be in too deep.

Finally, with the tools to hand, the Risk Manager is ready to start making difficult decisions. Her role is not simply to report the risk of doing nothing, but to manage the risk, by evaluating and choosing between different interventions. Each intervention, in changing the hazard outcome probabilities, the footprint function, or the loss operator, changes the EP curve. Ultimately, therefore, the Risk Manager will be faced with one diagram containing an EP curve for each intervention. As each intervention also has a financial and social cost, and also a probability of successful completion, the Risk Manager is not able to proceed on the basis of the EP curves alone, but must perform a very demanding synthesis of all of these aspects of the problem. It is not clear that anything other than very general guidelines can be given for such a challenging problem. But what is clear is that the Risk Manager will be well-served by a set of EP curves and probabilistic hazard maps that are transparently and defensibly derived. The purpose of this Chapter has been to make this derivation as transparent and defensible as possible.

Acknowledgements

I would like to thank Thea Hincks and Keith Beven, whose detailed and perceptive comments on previous drafts of this chapter lead to many improvements.

References

Cox, D.R., 2006, *Principles of Statistical Inference*, Cambridge University Press.
de Haan, L. and A. Ferreira, 2006, *Extreme Value Theory: An Introduction*, Springer.
Davison, A.C. and D.V. Hinkley, 1997, *Bootstrap Methods and their Application*, Cambridge University Press.
Davison, A.C., 2003, *Statistical Models*, Cambridge University Press.
DeGroot, M.H. and M.J. Schervish, 2002, *Probability and Statistics*, 3rd edn, Addison-Wesley.
Evans, M. and T. Swartz, 2000, *Approximating Integrals via Monte Carlo and Deterministic Methods*, Oxford University Press.
Gigerenzer, G, 2003, *Reckoning with Risk: Learning to Live with Uncertainty*, Penguin Books Ltd.
Grimmett, G.R. and D.R. Stirzaker, 2001, *Probability and Random Processes*, 3rd edn, Oxford University Press.
Hacking, I., 2001, *An Introduction to Probability and Inductive Logic*, Cambridge University Press.
Halpern, J.Y., 2003, *Reasoning about uncertainty*, MIT Press.
Hardy, A.R., W. Roelofs, A. Hart and D.A. Macleod, 2007, *Comparative review of risk terminology*, Technical Report, The Central Science Laboratory, doc. S12.4547739 (Nov).
Lindley, D.V., 1985, *Making Decisions*, 2nd edn, John Wiley & Sons.
Mason, B.G., D.M. Pyle and C. Oppenheimer, 2004, The size and frequency of the largest explosive eruptions on Earth, *Bulletin of Volcanology*, **66**, 735-748.
Norton, J.D., 2007, Probability Disassembled, *British Journal for the Philosophy of Science*, **58**, 141-171.
Paris, J.B., 1994, *The Uncertain Reasoner's Companion: A Mathematical Perspective*, Cambridge University Press.
Porter, J., 2009, *Lost in Translation: Managing the 'Risk' of Flooding – The Story of Consistency and Flood Risk Mapping*, unpublished conference paper, available at

http://www.kcl.ac.uk/sspp/departments/geography/research/hrg/researchsymposium/1aporter.pdf
R Development Core Team, 2004, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, http://www.R-project.org/
Rice, J.A., 1995, *Mathematical Statistics and Data Analysis*, 2nd edn, Duxbury Press.
Salt, J.D., 2008, The seven habits of highly defective simulation projects, *Journal of Simulation*, 2, 155-161.
Walley, P., 1991, *Statistical Reasoning with Imprecise Probabilities*, Chapman and Hall.
Wasserman, L., 2004, *All of Statistics: A Concise Course in Statistical Inference*, Springer.
Woo, G., 1999, *The Mathematics of Natural Catastrophes*, Imperial College Press.