Association of British Insurers

NON-MODELLED RISKS

A guide to more complete catastrophe risk assessment for (re)insurers

Written and edited by (re)insurance industry practitioners
Published by the Association of British Insurers
London, April 2014

Publisher **Association of British Insurers**

Editors David Clouston, Lloyd's
Tristan Garnons-Williams, Regulation Directorate, Association of British Insurers

Authors **(in alphabetical order of organisations)**

AIR Worldwide Ltd
Dr Milan Simic, Senior Vice President and Managing Director
Shane Latchman, Senior Research Associate

Amlin
David Singh, Group Underwriting Exposure Manager
JB Crozet, Head of Group Underwriting Modelling

AON Benfield
Catherine Tillyard, Catastrophe Management

Aviva
Chris Boss
Evonne Twite

Canopius Managing Agents
Marek Shafer, Head of Catastrophe Management

Catlin Group
Paul Kershaw, Group Aggregation Manager
Mojum Khan, Aggregate Modelling Manager

Guy Carpenter
Victoria Jenkins, Co-head of Analytics EMEA

Kiln Group
Lars Schmid, Model Validation Actuary

Liberty Syndicates at Lloyd's
Graham Clark

Lloyd's
Trevor Maynard

Munich Re
Hjortur Thrainsson
Alexander Allmann

PwC
Domenico del Re, Associate Director
Rebecca Wilkinson, Senior Actuarial Consultant, MA
Pau Ling Yap (non-modelled risks icons, designer)

QBE European Operations
Simon Webber, Head of Exposure & Catastrophe Management

Risk Dynamics SPRL
Dr Sebastian Rath, Principal Advisor

RMS Limited
Brian Owens, Senior Director
Dr Christine Ziehmann, Senior Director

RSA Insurance
Stephen Etheridge, Head of Strategy and Policy for Group Reinsurance, Actuary
Nathan Williams, Underwriting Director, Actuary

SCOR SE
Paul Nunn, Head of Cat Risk Modelling
Phil Holt, Head of Exposure Management, The Channel Syndicate at Lloyd's

Tokio Marine Technologies LLC
Dr Atul C. Khanduri, Executive Vice President
Dr Jianming Yin, Senior Vice President

Willis Re
Ian Cook
Matthew Eagle
Dr Rick Thomas

Zurich Insurance Group
Christof Grass-Fleury
Dr Matthew Jones

Table of Contents

Background

In November 2011, the Association of British Insurers (ABI) published 'Industry Good Practice for Catastrophe Modelling: a guide to managing catastrophe models as part of an Internal Model under Solvency II'. Written by a group of insurance and reinsurance professionals, it was well received by the industry, and achieved a significant worldwide audience.

In March 2013, the Financial Services Authority (FSA) – the predecessor organisation to the UK's Prudential Regulation Authority (PRA) – met in London with representatives from domestic and global insurers, reinsurers, reinsurance intermediaries, catastrophe model vendors and consultants.

The initial meeting formed part of a series of workshops from March to June 2013 in order to consider ways of encouraging the identification, quantification, monitoring, reporting and management of 'non-modelled' catastrophe risks. The PRA presented a high-level working framework based on its industry-wide understanding of the issue.

The group members acknowledged that industry-led guidance would be extremely helpful, and they agreed to collaborate on producing a document; drafting the technical content under the auspices of the ABI who agreed to publish it.

The PRA was kept informed throughout the composition stage, but was not responsible for defining or drafting the document. Accordingly, this document does not amount to PRA guidance and does not necessarily represent the PRA's views on non-modelled catastrophe risks.

Good practice, not current practice

In drafting this document, the authors have sought to focus on what constitutes 'industry good practice' for understanding non-modelled catastrophe risks.

They have, therefore, attempted to describe 'good' practice, rather than necessarily the 'current' practices in their own organisations or the wider market. It is in the spirit of sharing their own challenges and aspirations that these thoughts about industry good practice are presented here.

Who are the authors?

The authors all work in areas predominantly concerned with understanding natural catastrophe risk as it affects insurers and reinsurers, and representing that risk within companies' internal models.

A full list of the organisations represented by the authors and editors can be found at the front of this document.

The limitations of this document as regulatory guidance

Everything contained within this paper represents only the collective opinion of the authors. As a group of industry practitioners, they do not have any formal or regulatory status.

Therefore, readers should always bear in mind that the content reflects a best attempt at guidance, based on the authors' own expertise and experience. It does not in any way represent the views of any regulatory authority, including the PRA, or any other statutory or regulatory body.

Adopting the suggestions here will not mean, or imply, that an insurance or reinsurance company is necessarily meeting regulatory requirements. Wherever there is any apparent inconsistency between the text of this document and the relevant regulatory requirement, the latter must always be regarded as definitive.

In short, companies need to comply with all legal requirements placed upon them by regulators. Nothing in this document officially qualifies, limits or extends those requirements.

One size does not fit all

While the authors have done their best to suggest what constitutes 'industry good practice', they are aware that every regulated entity is different, and will have different needs.

Readers should, therefore, always test the relevance of these guidelines against the needs of their particular companies. In particular, what is 'proportionate and material' will vary from company to company, and even within different parts of each company's insurance or reinsurance portfolio. A reasonable requirement for one company may not be appropriate for another, and vice versa.

Notes on the text

Chapters and sections

Throughout this document, ‘chapter’ refers to the main topics listed in the table of contents; ‘section’ means the individual, numbered paragraphs within each chapter. So, for example, Chapter 2 is about Governance; Section 2.2.3 discusses expert judgement in that context.

All references to chapters and sections are hyperlinked throughout the paper for ease of navigation.

Acronyms

The authors have tried to keep the use of acronyms to a minimum. However, in a technical document, there is no avoiding them altogether. Each first use in the text is explained; thereafter, the acronym only is used. The most frequently used acronym is:

‘NMR’ FOR ‘NON-MODELLED CATASTROPHE RISK(S)’

Rather than spell out ‘non-modelled catastrophe risk(s)’ each time, the authors have used the acronym NMR throughout this document. Chapter 1 contains a definition of ‘non-modelled risks’.

Definitions

The following terms have the following specific meanings within this document.

CATASTROPHE MODEL

Probabilistic, exposure-based models of catastrophe risk affecting specific regions and perils (for example, US windstorm). Also known as catastrophe models, or Cat models for short, these are currently the most sophisticated and complex representations of natural physical catastrophe available to insurers and reinsurers. A catastrophe model comprises a catalogue of event sets, describing the multiple variations of a risk across a region and peril. It is referred to as a probabilistic model as it encapsulates the multiple variations of risk and the probability of occurrence.

There is a much more limited set of models for terrorism and infectious disease using related methods.

COMPANY

Insurance and reinsurance ventures come in many corporate forms, including mutuals, limited-liability companies, state-owned corporations and Lloyd’s managing agents. For the sake of uniformity, this document uses the term ‘company’ to represent all such entities.

DAMAGE FACTOR

The percentage of zonal aggregate sum insured that would be lost in a catastrophic event with a reasonably – but not extremely – low probability of occurring in any given year. The geographical zone may be any size (for example: postcode, CRESTA zone, or whole country).

‘Damage factor’ is now sometimes used synonymously with Probable Maximum Loss, which is an estimate of ground-up loss relative to sum insured (see below). The authors use the term ‘damage factor’ throughout this document.

LOSS FACTOR

The percentage of premium income – rather than exposure – that would be lost in a catastrophic event.

PML

Probable Maximum Loss originally applied to individual locations within a Fire & Allied Perils policy schedule – for example, a factory complex – where the total sum insured could not reasonably be thought at risk from one serious fire or explosion. The PML was therefore the largest loss that could be sustained from a serious event affecting one location, but generally not taking into account natural catastrophes.

Over time, PML has acquired the looser meaning of a percentage of zonal aggregate that would be at risk from catastrophe. However, the authors use the term ‘damage factor’ (see above) throughout the paper to express this concept.

1. What are 'non-modelled' risks?

1.1 Introduction

Following the major natural catastrophe events of 2010 and 2011, the term 'non-modelled risks' entered the insurance industry's lexicon as short-hand for limitations in the suite of tools commonly used for catastrophe risk management.

In particular, the devastating Thailand floods in summer 2011 emphasised the need to recognise sources of catastrophe risk not represented in widely-available vendor models.

It is important to stress that such risks typically were modelled in some form, where companies considered them to constitute a material source of risk. However, the modelling often used relatively high-level techniques that were not necessarily exposure-based. Losses to some companies called into question the effectiveness of these techniques, given the materiality – as it proved to be in practice – of the exposure.

Similarly, while catastrophe models are well established for Japanese earthquake scenarios, none of the widely-available vendor models in March 2011 represented the potential for concurrent tsunami damage and losses.

Modelling these – or other – regions may not necessarily affect capital calculations for companies with more material risks elsewhere, but they do underscore that NMR can lead to earnings' volatility, and can start to become more relevant to capital as exposures grow.

1.2 Definition

The authors aim to provide practical guidance for closing the 'non-modelled' gaps and moving towards global completeness of natural hazard modelling. For the purposes of this document, therefore, NMR is defined as:

'Any potential source of non-life insurance loss that may arise as a result of catastrophe events, but which is not explicitly covered by a company's use of existing catastrophe models.'

'Non-modelled catastrophe risk' means different things to different organisations, and there is no market-wide understanding of what constitutes 'non-modelled' perils.

It is also important to stress that internal models used for capital and solvency purposes should – at least implicitly – take into account all material sources of catastrophe loss.

This paper, however, aims to support companies in improving the explicit identification and modelling of catastrophe risks within their risk management and capital modelling frameworks. This includes methods of representing such risks in their capital models (see Chapter 7).

1.2.1 Non-modelled vs. uncertainty in relation to modelled losses

There are considerable uncertainties implicit in the outputs of vendor catastrophe models, even for theoretically well-understood regions and perils. For the purpose of this paper, however, these uncertainties are out of scope (see Section 1.3.3).

1.2.2 Dynamic modelling landscape

Catastrophe models are continually improving, both in terms of the breadth of region/perils covered, and the sophistication of their coverage. Every year, as lessons are learned and investment in emerging/developing economies presents new concentrations of insured risk, more gaps in the models are closed. Any attempt at an exhaustive catalogue of NMR would be rendered almost immediately obsolete.

While the examples used may age over time, the authors hope that the approaches suggested will remain applicable even after today's particular gaps have been filled.

1.3 Document scope

To make this a practical document both to write and read, the authors have explicitly limited the classes and lines of business to be covered within the paper.

1.3.1 'In scope' lines of business

The following lines of business are within the scope of this paper:

- property
- auto
- engineering/construction
- energy
- marine (including fine art/specie)
- contingency
- aviation
- accident & health/WCA
- terrorism (included due to the similarities of techniques that lend themselves to assessing fixed-asset concentrations in natural disasters)

1.3.2 'Out of scope' lines of business

The following lines of business are not within the scope of this paper:

- **agriculture**. While correlated in some flood/hail/freeze events, agriculture risks require somewhat different techniques to address.
- **liability classes**. Certain liability classes have generated losses in past catastrophe events (for example, nursing home liability in Hurricane Katrina). However, on materiality grounds, the focus will be on the 'P' rather than the 'C' of the Property and Casualty (P&C) universe. Likewise, systemic liability risks are also considered out of scope (for example: asbestos, electro-magnetic fields, etc.), although mould and environmental liability associated directly with natural hazards events are considered.
- **life insurance**. Life insurers who have some exposure to natural catastrophes (for example, the 2004 Indian Ocean tsunami) may find the paper helpful. However, Life data and coverage knowledge fall outside the expertise of the authors, who are drawn from P&C companies.

1.3.3 Model uncertainty

- as noted in Section 1.2.1, limitations and uncertainties in the catastrophe models are considered out of scope. For the purposes of this paper, they fall within the area of model validation rather than NMR[1]. Please refer to Chapter 7 of the ABI's 2011 publication 'Industry Good Practice for Catastrophe Modelling'.

1 There have been notable examples of losses that have 'surprised' the industry, despite their inclusion in catastrophe models available at the time. For example, the magnitude of the 2011 Tohoku earthquake event was outside the range expected from the Nankai Trough. Both the 2010 Christchurch and 1994 Northridge earthquakes involved faults not explicitly identified in otherwise seismically active areas. In each case, however, the region and peril involved was included in contemporary catastrophe models, albeit in a way which did not fully represent the risk. Therefore, the sources of loss for these events are considered to be modelled rather than non-modelled.

1.4 Types of NMR

Sources of NMR loss can be considered in four main categories. This categorisation provides a useful framework for considering the most significant challenges addressed in this paper.

1.4.1 Regions and perils not covered by catastrophe models

Throughout the text, the icon opposite will indicate regions and/or perils not covered by an existing available catastrophe model, or a model that a company chooses not to licence.

The real world constantly highlights material catastrophe risks not incorporated into available analytical tools. Examples include the World Trade Centre loss of 9/11, and the Thailand floods in 2011. In both cases, there has been extensive post-event activity to fill particular gaps in the modelled landscape.

1.4.2 Secondary perils and effects not covered by catastrophe models

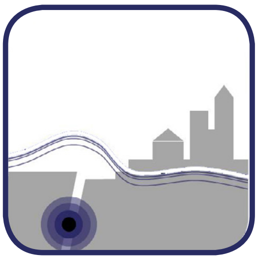

The use of the icon opposite indicates perils whose primary event characteristics are captured in a catastrophe model, but whose losses from resultant or secondary perils are not represented. An example might be earthquake shaking (included in the catastrophe model) and subsequent tsunami (not included). This may vary from one model vendor and region to another at any point in time.

While some of the secondary perils listed below may be modelled in some vendor – or other – models in certain regions, examples of secondary perils may include:

- storm surge
- fire following earthquake/terrorism/flood
- tsunami
- liquefaction
- landslide/mudslide
- looting
- demand surge or post-loss amplification
- loss adjustment expenses
- regulatory intervention
- other consequential losses

Secondary perils can cause very material losses. From our 'top 12' catastrophe events (see Section 1.5), we see significant non-modelled secondary impact in:

- Tohoku earthquake tsunami and the Fukushima nuclear crisis
- Hurricane Katrina levee failure
- Hurricane Ike inland wind/flood
- Superstorm Sandy fire following flood, looting, pollution

Correlation between 'separate' events (for example: aftershocks, clustering, etc.) can also be considered as secondary effects.

1.4.3 Classes and lines of business not covered by catastrophe models

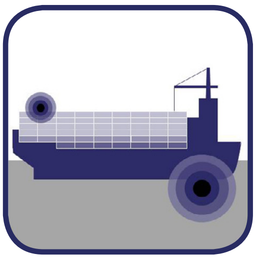

While many companies routinely use catastrophe models to capture and model property classes, this does not always extend to other classes which have the potential to sustain losses in an event. The icon opposite indicates these throughout the paper.

Examples may include:

- workers compensation or personal accident
- energy/power risks (offshore platforms, wind farms)
- inland marine
- engineering and infrastructure (motorways, bridges, special structures)

Non-static risks such as goods in transit, hull, cargo, art exhibitions, and yachts are a particular challenge. Superstorm Sandy once again exemplifies this challenge with significant losses for cargo, fine art and recreational marine classes.

1.4.4 Coverages not considered by catastrophe models

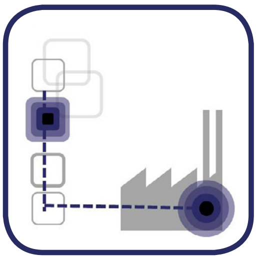

Most catastrophe models routinely cover physical damage and business interruption. However, additional coverages and/or sub-terms and conditions may not be explicitly modelled. These are indicated by the icon opposite.

For residential risks, this may include freezer contents or additional living expenses. Commercial and/or industrial examples include pollution, debris removal and machinery breakdown.

Offshore energy risks have especially complex coverages including, among others: Control of Well; Removal of Wreck; Re-drill; and more. Again, some catastrophe models licenced by a company may represent some or all of these; others may not.

A topical example in both the Tohoku earthquake and Thailand floods is the potential for contingent business interruption (CBI) losses, as suppliers affected by the disasters are unable to meet contracts to their customers around the world.

1.5 Impact of NMR

The diagram below represents the twelve largest catastrophic events (by insured loss) in history.

The tree map in Figure 1 shows data in nested rectangles. Each event has been classified in two ways:

- red/amber/green (RAG) indicates the approximate relative contribution of NMR to the insured loss. As noted in Section 1.2, the definition of NMR is subjective and based on different modelling capabilities at the time. The RAG rating is an estimated overall view of the industry experience, based on a consensus of the authors.
- the size of area covered by each event represents the relative industry-loss size.

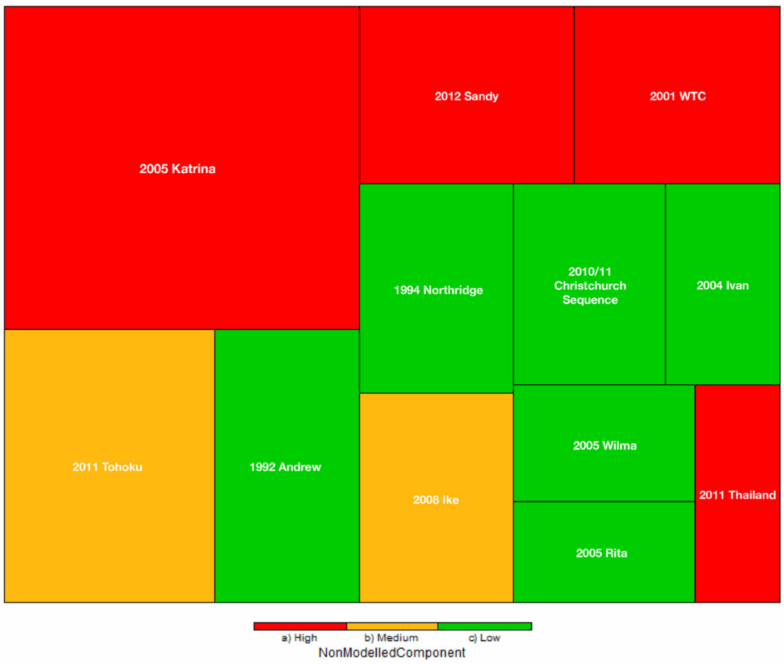

Figure 1: Non-modelled risks tree map

Table 1 below, classifies the non-modelled aspects of each loss.

EVENT	CATEGORIES	DESCRIPTION
2005 Katrina	Secondary peril	Levee failure
		Extreme loss-amplification of 'super cat'
	Classes not incl.	Non-property Losses from a cat event
		Policy 'Leakage' (e.g. flood losses recovered from a wind policy)
Tohuku		Tsunami
		Fukushima (Nuclear Power)
	Coverages not incl.	Contingent Business Interruption (CBI)
Sandy		Government Intervention in application of deductibles
		Non-modelled lines of business (auto, marine)
		Pollution
		Fire following flood, etc.
WTC	Region/Peril	Terror

EVENT	CATEGORIES	DESCRIPTION
Andrew	N/A	Windstorm only
Northridge	N/A	Earthquake only
Ike		Secondary effects in Mid West & North East
		Inland Flooding
Christchurch Sequence		Clustering / Aftershocks
		Liquefaction / Landslide
Ivan		Subsea pipelines Miss. Delta
Rita	N/A	Windstorm only
Wilma	N/A	Windstorm only
Thailand		Region/peril not modelled at the time of the event
		Flooding in industrial parks exacerbated by the need to defend Bangkok
	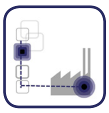	CBI
	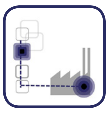	'Interests abroad' coverage

Table 1: Non-modelled risks

2. Governance

2.1 Introduction

As with all aspects of risk management, explicit corporate processes need to be in place to identify, quantify and report NMR. Management of NMR should be properly embedded within a company's culture.

One important caveat must be stated: processes to identify new NMR can be subject to two great enemies of effective risk management: familiarity and routine. Emerging risks may be habitually ignored, since they tend to lie – almost by definition – outside existing processes. It is extremely important that the management cycle specifically includes structured, open 'thinking time' to consider what is not being seen.

The governance of NMR may be categorised into three main areas:

- organisation and methods. Who is responsible? What are the roles? Are methods and tools different to those used for the modelled risks?
- cyclicality and triggers. Where does identification stop and quantification start? When is the implementation of the modelling approach reviewed?
- communication. How is comfort in the modelled output communicated? How can management information be designed to help broader aspects of the business?

This Chapter examines all three areas, and briefly considers means of embedding a 'NMR culture' within a company.

2.2 Organisation and methods

In deciding which corporate function/area of a business should manage NMR, companies should consider what skill sets are relevant to the task and who has them within the organisation. This Section is intended to help companies make their own decisions about the areas of expertise required to manage and quantify NMR in their organisation.

2.2.1 Tailoring methods to NMR

Relevant skill sets could include: exposure management, development of loss models, model evaluation and analysis of claims history. One essential skill set is the understanding of currently-licenced catastrophe models and their limitations. Those responsible for modelling perils using explicit methodologies, such as probabilistic catastrophe models, will have expertise or strong familiarity with the components of catastrophe modelling that give them the tools to take a structured approach necessary to manage NMR.

Solvency II requirements have clarified and enhanced understanding of the limitations and credible use of probabalistic models, which can help companies to better understand how they manage NMR.

Companies should consider what tools are available in the market, from probabilistic vendor models to basic data layers, all of which are constantly evolving. As a result, the boundaries between modelling approaches should be reviewed on a systematic, regular basis. One group of experts in a company should be tasked with considering whether appropriate methods are still being used in each case.

These principles suggest it might be beneficial for a company to establish a single point of responsibility covering all catastrophe risks. Each organisation can decide where the point of responsibility is, and who owns it; be it the risk, actuarial, or underwriting departments, or a combination thereof.

2.2.2 Materiality

Uncertainty in the quantification of NMR will tend to be higher than for explicitly-modelled catastrophe risk. Uncertainty makes the assessment of materiality more complex for NMR. This may require adjustments to how a company's materiality thresholds are defined, and how calculations are performed to understand if the thresholds have been exceeded. For example, a NMR that may appear non-material on average should potentially be included in risk assessments if the uncertainty around its quantification is considered to be high.

2.2.3 Expert judgement and peer review

NMRs are likely to be such because there is little hard exposure and/or claims data to support the modelling. The use of expert judgement may, therefore, be more material. Governance relating to the extent of expert judgement – and the peer review put in place – should be considered throughout the process, particularly relating to validation activities.

2.3 Cyclicality and triggers: risk management cycle

The process of dealing with NMR should follow a standard risk management control cycle, as illustrated below in Figure 2.

Figure 2: Risk cycle chart

2.3.1 Establishing context

The first stage in the process should be to establish the context in which the risks should be measured. For example, if the overall objective is to achieve a more accurate capital number for solvency purposes, looking at tail events will be much more important than if the purpose is for the pricing of risk (see Section 4.3.1 for more on this topic).

Context will also inform whether the risks that need to be identified and managed are those to which the company is currently exposed, or those to which it may be exposed in future. For example, a new risk could result from a company writing new classes of business or building a portfolio in a new market, or a legislative change that compels insurers to provide coverage that was not previously provided.

2.3.2 Identification

Once the context is established, risks need to be identified. Chapter 3 discusses different approaches to identification. As noted in the introduction to this Chapter, constant vigilance is required as part of the identification process and should be built in to the formal risk management framework.

2.3.3 Analysis

Though analysis and evaluation can constitute two distinct activities (potentially involving different individuals or corporate functions), a project may switch repeatedly between these two activities during the first time a company embarks on the risk management cycle. Approximate, and therefore uncertain, analyses may precede an initial evaluation. This could result in the need to refine and deepen the analysis until the quantification can support a decision on materiality.

2.3.4 Evaluation of risks and materiality

Once the material risks have been sufficiently analysed, the evaluation stage should include reporting to senior management, as well as considering and implementing a plan to address risk. This could include mitigating the risk through reinsurance, putting in place a monitoring regime, or embarking on more detailed risk analysis. Those companies in compliance with Solvency II's model validation requirements will also need to perform a regular validation of risk modelling. As part of this monitoring process, the steps in the risk cycle should be documented and refreshed as appropriate.

A company should also consider which parts of the cycle are repeated as part of the regular validation routine, and which may be triggered by business cycles, management actions or actual catastrophe events.

The types of triggers for a re-analysis or new identification of NMR may include:

- a catastrophe event that has identified a cause of loss not previously considered. For example, in 2010 and 2011 the insurance industry was affected by catastrophe events from 'non-peak' regions, and the World Trade Centre loss of 9/11 triggered a re-think about dependencies between classes of business.
- newly available modelling techniques and/or data sets
- scientific and/or legislative changes
- regulatory scrutiny and/or more stringent capital requirements
- new strategic focus and/or business plans

2.4 Communication

This Section focuses on the specific communication issues related to NMR, rather than catastrophe risk in general.

Communication objectives specific to NMR could include:

- interface between the stakeholders in the technical evaluation (for example, underwriting and capital modelling teams)
- explaining prudence and residual uncertainty in the modelling
- demonstrating to senior management the costs and benefits of additional investment in processes and modelling

2.4.1 Interface between teams

As highlighted above, clear points of responsibility need to be in place to manage catastrophe risks, with particular attention paid to the boundaries between different approaches to dealing with NMR.

For companies likely to be operating under Solvency II, the capital model should make some allowance for all identified exposures, and capital teams will have generally dealt with these allowances. When companies identify material relating to NMR, it can result in responsibility shifting to other teams in the organisations, and the cyclical nature of the identification process means that ownership of modelling NMR could be challenged and boundaries shifted.

A common taxonomy should exist inside the company to avoid misunderstanding between teams. Training may be required to ensure the expertise of any one team is shared. For example, the company should understand how potential correlations of losses between classes from a physical event are fed through the capital models and any related limitations. The team managing the quantification of this risk should be informed on how to structure the model and its output.

2.4.2 Prudence and residual uncertainty

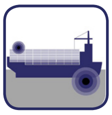

At every stage of the process, it is important to highlight what is and what is not included, and what residual gaps and uncertainties remain.

In particular, when risks are quantified using a streamlined approach, the modelling function will normally have built a view on the coverage achieved in the approach, and the level of comfort in the analysis. Catastrophe models especially can suggest spurious accuracy and there may be a risk of issues being missed if the risk is considered 'modelled'.

2.5 Embedding NMR culture

If culture can be understood as "how we do things around here", then only when the management of NMR has been truly embedded within a company's culture can the benefits be fully realised.

There are a number of factors required to effectively embed NMR within a company:

- senior management buy-in. This can be attained by demonstrating the future benefits of embedding a NMR culture. The benefits of improved management information to the business may include:
  - better exposure management from the knowledge of risks that are affected by regions and perils, classes, coverages and secondary perils and effects
  - better data capture to facilitate future NMR assessment and quantification (for example, motor flood and wind damage is often recorded just as 'property damage', which can impact the ability to assess past losses from these perils for this class of business)
  - enhanced reinsurance purchasing
  - more accurate risk pricing
  - better management of underwriting volatility and trends
  - supporting the growth strategy of the company
- appropriate resources assigned to the task
- incentives to support the company's NMR strategy
- progress and results being measured and communicated within the organisation
- good documentation practices which apply just as much to NMR as other risks and, in particular, where bespoke methodologies are developed internally

3. Identification

3.1 Introduction

Chapter 2 discussed the necessity of having a clear, structured process for managing NMR, with identification being the first step. There should be a systematic, documented process for identifying NMR that employs appropriate resources and techniques. Once identified, the materiality of the NMR should be assessed. The process itself should be regularly reviewed in accordance with company policy in order to guard against familiarity and routine.

There are – very broadly – three methods for identifying NMR:

- exposure-based: analysing current or planned exposures
- claims-based: analysing past claims
- expert judgement: gathering expert opinions and analysis

Each method has its challenges, and most require significant resource. However, systematic assessment of catastrophe risk that is not currently modelled will almost certainly include all three methods to some extent. Each is discussed in turn below. The chapter concludes with some practical suggestions for maintaining vigilance.

Figure 3 is an example of the sort of NMR identification and assessment process that could be employed. At each of the decision points, all three of the above methods would be employed to understand if a NMR existed and, if so, its materiality.

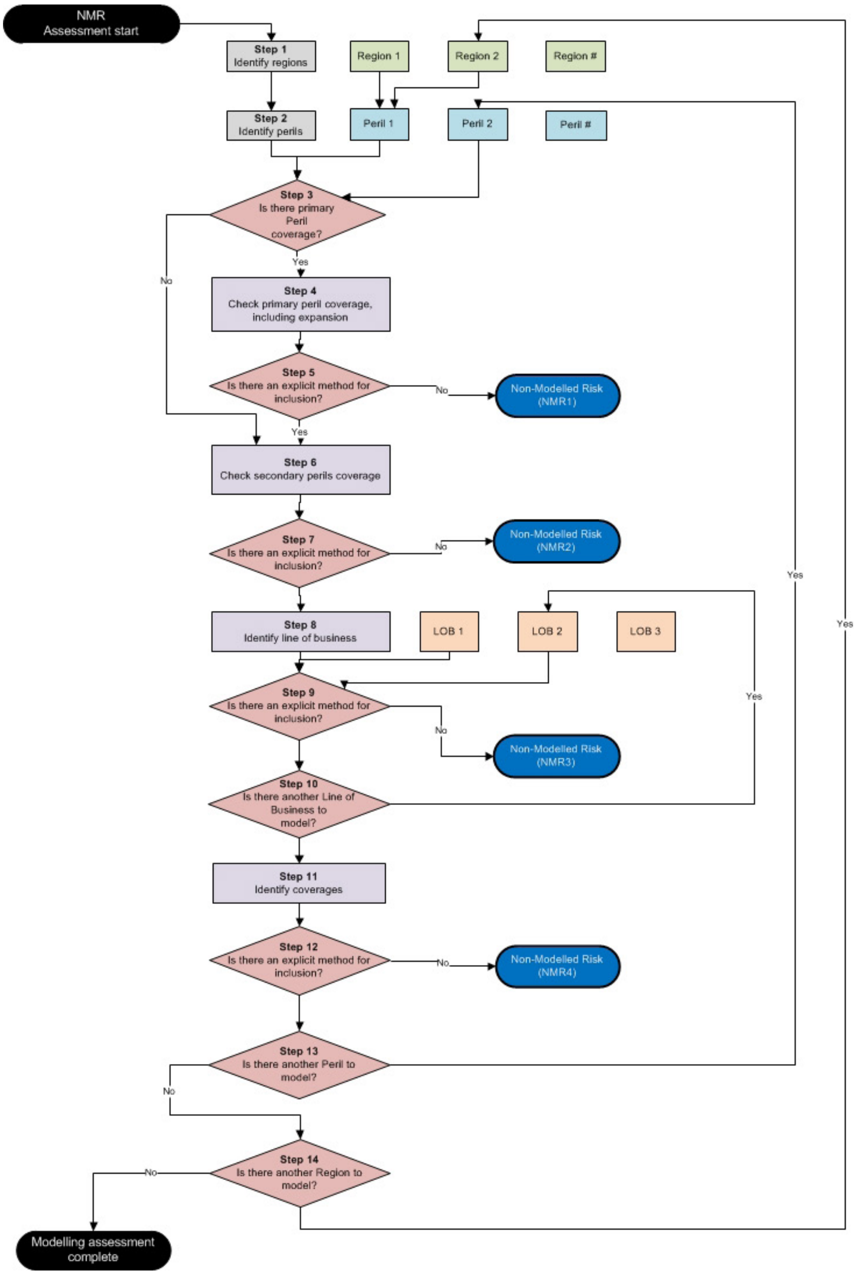

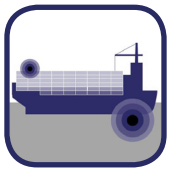

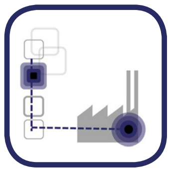

Figure 3: Process map

3.1.1 A note about materiality

It is impossible to overstate the importance of taking into account the concept of 'materiality' when considering NMR. All other things being equal, the most time and effort should be spent considering the most material non-modelled exposures at risk. Identifying these should always be the starting-point of any investigation into NMR, or an assessment of their materiality.

Companies likely to be operating under Solvency II should adopt a working internal definition of materiality. The following sections of this paper describe various methodologies to estimate levels of catastrophe risk not captured by more explicit methodologies. The appropriateness (or otherwise) of those methodologies will depend on the materiality of the risk for the company in question at the time of calculation.

A company should ask itself:

- on what basis has materiality been defined? Is it one or more of a combination of, for example, capital held, share price, market share, risk appetite, reinsurance purchase and/or spend, or credit ratings?
- how will the results be used and will this potentially affect the materiality threshold?
- how might the uncertainty and volatility of NMR affect the materiality thresholds?
- for capital-relevant return periods, what is the materiality of catastrophe risk within the company's overall risk?
- what is the relative materiality of different regions/perils?
- how does the company monitor materiality of catastrophe risk in future?
- do the materiality thresholds consider mitigation arrangements, such as reinsurance?
- has due consideration been given to the correlation between region and perils, classes of business, coverages, and modelling methodologies?

3.2 Exposure analysis

Exposure analysis starts as a top-down process. Where dictated by materiality, a bottom-up analysis will add granularity and perspective.

3.2.1 Understanding current ground-up exposure to primary perils

Starting with on-risk policies, the company should draw up a list of all the countries where it provides coverage, followed by a systematic assessment of the materiality of its exposure in each.

Companies should be wary of the distinction between the origin/domicile of the risk and where the actual exposure is, which may be quite different. For example, the head office of an insured party may be in the United States, which will be the domicile, but the schedule of locations may be worldwide, including material exposures in other countries.

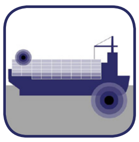

For territories or regions where exposure is sufficiently material, this principle should be followed for each line of business. A company with a property account will need to look at personal lines, commercial, industrial risks and engineering exposures; for example, Contractors All Risks and Construction All Risks (CAR). This should allow the company to create a detailed picture of the lines of business covered, and any adjustments that may be appropriate, for example, factoring in the projected exposure to the end of the contract term for CAR.

Primary catastrophe perils should be evaluated against each country or region. Expert opinion may be required to compile a comprehensive peril list. Where there is insufficient policy-level data to identify regional exposures in a country, the company may wish to assume itself exposed to all perils that may affect any part of the whole country. For example, a policy located in Melbourne, Australia will have little exposure to cyclones, but a policy identified only as 'Australia' should be assumed to be significantly at risk from that peril.

Once again, the importance of materiality cannot be overemphasised. The materiality of exposure in a country will drive the depth and breadth of all subsequent investigation. For example, a company may appropriately decide that it is easier to assume a total loss of all exposed limits for all possible identified perils in a country, rather than try to model them.

3.2.2 Understanding catastrophe model coverage and secondary perils

The next step in the process should be a review of the perils covered by licensed catastrophe models or existing modelling processes within the company. It is unlikely that any vendor model suite will encompass all the natural perils to which the company is exposed. It is essential, therefore, that a 'gap analysis' of the missing perils is undertaken. If the vendor tool covers earthquake, windstorm and river flood then all other perils such as freeze and hail may not be modelled.

It is crucial that a company has an appropriate level of knowledge – taking into account materiality of exposure – about vendor models' limitations as assessed during model validation. This can only be achieved by actively tracking vendor model documentation, asking questions and having model-specific discussions with vendors or reinsurance brokers. No model covers every eventuality, and the company must maintain appropriate awareness of the limitations for the four categorisations of NMR discussed in Section 1.4.

For example, a company may note that earthquake is covered for a specific territory in its property account. However, the prospect of fire following earthquake should be addressed. Is 'fire-following' a significant risk to the company? If yes, does the model capture this exposure and assess it? Does the tool adequately cater for this risk? What does the vendor's model documentation say?

Flood is another example. It is expensive to model and requires very detailed exposure data. What kind of flood is the company exposed to in the region? Is it surge, fluvial, pluvial, groundwater or surface water run-off that creates loss, or possibly all of them combined?

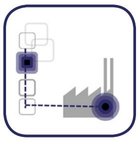

Where indicated by materiality, it is essential to scrutinise the classes of business written against the data input for the vendor model. What lines of business provide exposure data to the vendor tool? For example, if a company writes contractor risks, does the aggregate data from this class get coded in the model?

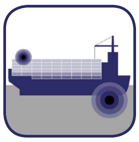

For what lines of business does the vendor tool have vulnerability curves? Models do not cater for all lines of business written and it is important to identify the gaps. For example, personal accident policies with UK territorial scope may not be represented by a specific vulnerability curve in the catastrophe model. Could this lead to material underestimation of the potential loss from, say, windstorm?

Having evaluated the primary perils, companies should then analyse secondary perils and effects. These perils, whilst not necessarily being the major source of loss, are important as they can produce significant additional losses to the primary peril modelled. As a result, they increase the severity of the worst-case scenario. Some examples are outlined in Table 2.

Earthquake induced	Tropical cyclone induced
Fire following	Storm surge
Sprinkler leakage	Inland flooding
Tsunami	Tornadoes
Liquefaction	Dam/Levee failure
Landslide	Debris removal
Flood following	Subsidence
Contamination clean-up	Contamination clean-up
Debris removal	Fire Following
Subsidence	
Dam failure	

Table 2: Primary/secondary perils

3.2.3 Understanding planned changes of exposures

Insurance portfolios change with time due to factors such as business strategy, legislation and availability of capital. Many companies plan to grow, and potential areas of growth are often in jurisdictions where the modelling is either untested or emerging. An analysis of current exposure or past claims experience may be insufficient, in their own right, to identify NMR.

To ensure that the risk identification process is proactive and not exclusively reactive, companies should consider any changes to the underlying portfolio. Regular communication about the company's growth plans should trigger research into potential hazards in growth areas, along with insurance market penetration. If significant growth is planned in an area of high hazard and low historic insurance penetration, companies should carry out further analysis. The same principle applies to new classes of business.

Understanding growth targets will also ensure that calibration of modelled catastrophe risk is undertaken within the appropriate parameters. Companies should ensure that data fed into models takes into account planned growth in the portfolio to avoid the risk of underestimation.

3.2.4 Understanding policy wordings

Policy wordings differ between companies and can be a useful tool for exposure management, but if they are not understood properly, they can also be a potential source of NMR. It is important that companies have a process in place to examine sources of loss within policy wordings and compare these against losses covered by modelling methodologies. This is best done in conjunction with underwriters and claims handlers, as they should be aware of the purpose and intention of the policy, and any potential consequences and ambiguities of certain clauses that may arise. Identifying gaps between policy coverages and the actual data captured by IT systems is of particular importance.

This process should also focus on establishing whether there are significant 'incidental' non-modelled additional exposures being assumed through lax exclusions, or broad policy wordings for lines of business that appear well-defined. There is a tendency for coverage to broaden in a softening market, so it is important to coordinate with underwriters to monitor these changes as an underwriting cycle progresses.

3.2.5 Understanding contagion and clash

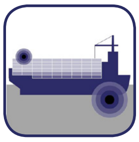

It is always important to be mindful that natural events do not occur in isolation, and that there is frequently interdependence between events. Companies should, therefore, identify any correlating factors between different classes of business. For example, the risk of flood is often considered in terms of its impact on the property account; however, 'motor own damage' may be a significant contributor to loss.

3.3 Claims analysis

The most straightforward and easily-available source of information to identify NMR is often claims experience. This can either be the experience of the whole industry or a company's own claims. Both sources are necessary to ensure the widest possible view of NMR.

3.3.1 Industry claims experience

Following most significant catastrophe events, a great deal of information is published that can help identify a company's sources of loss if not explicitly modelled. These sources of information might range from media reports of events (often useful for scene-setting) to detailed post-event reconnaissance reports by catastrophe model vendors. Recognised literature examples include reports from PCS, PERILS, AXCO, reinsurers and brokers, and academic publications such as the Earthquake Engineering Field Investigation Team (EEFIT).

Education about past claims can serve to unearth NMR that might otherwise have gone unnoticed. For example:

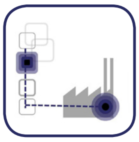

- what perils have caused significant losses in the past and does the company explicitly model these perils, coverages and classes?
- are there procedures in place to address missing areas?
- what about ancillary sources of loss, such as contingent business interruption (CBI)?
- the industry has experienced catastrophe events that have occurred in clusters, such as the European windstorms in 1990 and 1999. This clustering effect can also occur for other perils such as flood. Has the materiality of this effect been considered?

As always, companies should be aware of materiality when analysing industry information. For example, a particular NMR such as tsunami losses following an earthquake may be identified in an industry event report that is not relevant for the company because tsunami losses are excluded in the policy wording. In addition, industry experience has a wide scope and may, therefore, be too vague or general to use as the sole basis for numeric assessment.

3.3.2 Own company claims experience

The advantage a company has in analysing its own data is that it is directly relevant, and is likely to be available in detailed and granular form. The loss details are easily accessible, and can be accurately matched to policy data.

Of course, the scope of the lessons to be learnt can sometimes be more limited than from industry experience, depending upon the extent to which the company is exposed to an event. The stability of a company's portfolio can be another limitation. For example, new business is unlikely to have produced the breadth of claims experience that mature business would have produced.

Nonetheless, own claims experience is a crucial source of information. Examples of the type of information that can be derived from own claims experience include:

- claims without apparent corresponding exposures. Is there a non-reported exposure issue?
- assessment of the appropriate sum insured amounts, per the loss adjustor's assessment. Is there an insured-to-value (under-insurance) issue?
- analysis of the amount of Allocated Loss Adjustment Expense (ALAE) as a proportion of the indemnity amount, compared to the amount currently factored into the modelling

- review of the secondary causes and effects of loss emerging in the claims record that are not considered in modelling. For example, fire-following-earthquake, tsunami, or looting.

In the aftermath of a catastrophe event, a company will often begin a loss quantification process to review the adequacy and completeness of the modelling. The loss-modelling team will refer to model documentation and its own model validation assessment to ascertain the 'modelled' aspect of loss.

Forensic analysis of individual events can then be used to help identify NMR. Such analysis consists of reconciling the modelled estimate with the actual loss and following any significant discrepancy to its root cause. Typically, the sources of discrepancies may include:

- exposure information
- the event as modelled
- the modelling process itself
- non-modelled elements

Many discrepancies arise simply from the nature of probabilistic catastrophe models, which, necessarily, employ the rule of large numbers. This provides good insight into the overall loss, but may display significant variances from the perspective of individual risks.

The 'waterfall' chart below in Figure 4 shows how forensic analyses can help identify the non-modelled elements within the claims from a particular event. In Figure 4, RCV relates to 'replacement cost value' (i.e if a content item is lost, the insurer will pay the full value to replace the item like-for-like) and ACV to 'actual cash value' (i.e if a content item is lost, the insurer will pay the depreciated value of the item in question).

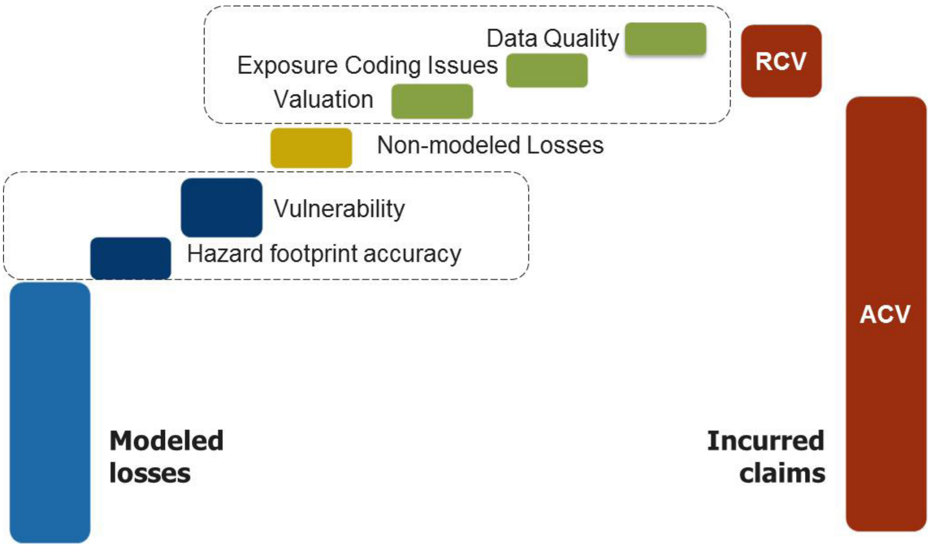

Figure 4: Waterfall chart

3.4 Expert opinion

Although exposure and claims analyses are essential techniques for properly identifying NMR, these are likely to be incomplete without the structured inclusion of expert opinion. We have noted previously that expert input may be required to identify primary hazards in different countries.

Expertise comes in many forms, including in-house subject matter experts, academics, external consultants, publications and conferences. However accessed, it is crucial that the use of experts is structured enough for a tangible outcome, yet open enough that all involved are encouraged to think of risks that may not happen on a regular basis, or may never have happened at all in recorded history.

This is not an easy balance to maintain.

Establishing an emerging risks committee and using techniques like ‘brain storming’ to discuss and think through a tangible scenario like a realistic disaster scenario, for example, - or set of scenarios - can help generate ideas for possible NMR.

Below are some suggestions for the type of expertise that can be helpful in terms of NMR identification.

3.4.1 Underwriting experience

Underwriters are closest to the original exposure and risk, so it is essential to harness their knowledge, which is a good source of insight into the likelihood of certain NMR occurring. A workshop in which underwriters use their experience to identify risks that could otherwise be easily overlooked is likely to provide a probing insight into the types of risks involved.

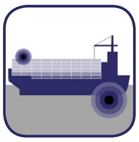

Again, a structured approach is likely to provide the best results. An assessment involving regions and lines of business should be reviewed. This assessment should be cross-checked against that used in understanding the ground-up exposure. It is important that regional exposure is fully documented and, having established this, underwriters should use their knowledge to assess the perils coverage provided. Again, this should be cross-checked against models, such as those from vendors.

3.4.2 Claims experts

In much the same way that underwriters can be a valuable source of information, so too can claims staff. The information obtained from handling claims provides a particular focus on risks that may otherwise fall under the radar. Again, a formal structure should be set up and implemented to ensure feedback from these experts is recorded, and actions taken, on a defined basis; for example, semi-annually in order to evaluate the information obtained from claims-handlers.

3.4.3 Risk engineering

Risk engineers and loss adjusters who visit insured locations before any claims are paid can be an excellent source of information. As part of their work, they routinely assess potential risks to sites and will be available to advise on items of risk that are potentially non-modelled. For example, a risk engineer might identify that the site is next to a steep hill and is therefore at risk from landslide. Conversely, the risk engineer may also identify that the site has substantial local flood defences and so can advise on a reduced flood exposure.

3.4.4 Catastrophe modelling vendors

Catastrophe model vendors can provide a good – and often unique – perspective on how to think about perils that are not yet modelled. Part of their expertise lies in the structured thinking that should underpin a company’s processes in relation to NMR. For example, asking questions such as ‘what are the essential characteristics of this risk that will enable us to represent it in a model?’

3.4.5 Local knowledge and expertise

Local underwriting teams and market practitioners often have detailed knowledge of their markets. This is particularly useful when the market is growing and changing, as they should be able to highlight trends in coverage or exposure growth that may be very difficult to identify from central locations or via top-down analysis.

The caveat to the above is that there is a tendency for stable markets to ignore potential perils that haven’t yet produced insured losses. This can be the result of luck, rather than being due to the absence of hazard. This implies that local opinions need to be balanced against an objective consideration of both feasible peril events, and the validity of the observation period that underlies the local knowledge. The financial implications of local arrangements should be evaluated and included within the modelling as these will be correlated with event losses.

3.4.6 Participation in expert groups

Companies that participate in expert groups within their industry are likely to be kept abreast of new developments and thinking. These can be an important source of information, provided that both attendance and subsequent sharing of information are properly structured.

3.4.7 Hazard information

There are a significant number of hazard maps available, either in the public domain or on a subscription basis from local authorities, specialist risk organisations and reinsurers. Some of these maps provide a global overview, while others are focused on details pertaining to a specific region/peril. These maps may be used to identify natural perils/regions not already covered by more sophisticated models.

For instance, Cartorisque, published by the French Government, provides very granular maps for flood, avalanche and earthquake risk in France (Figure 5). Nathan World Map of Natural Hazards, published by Munich Re, covers a wide range of perils worldwide (Figure 6).

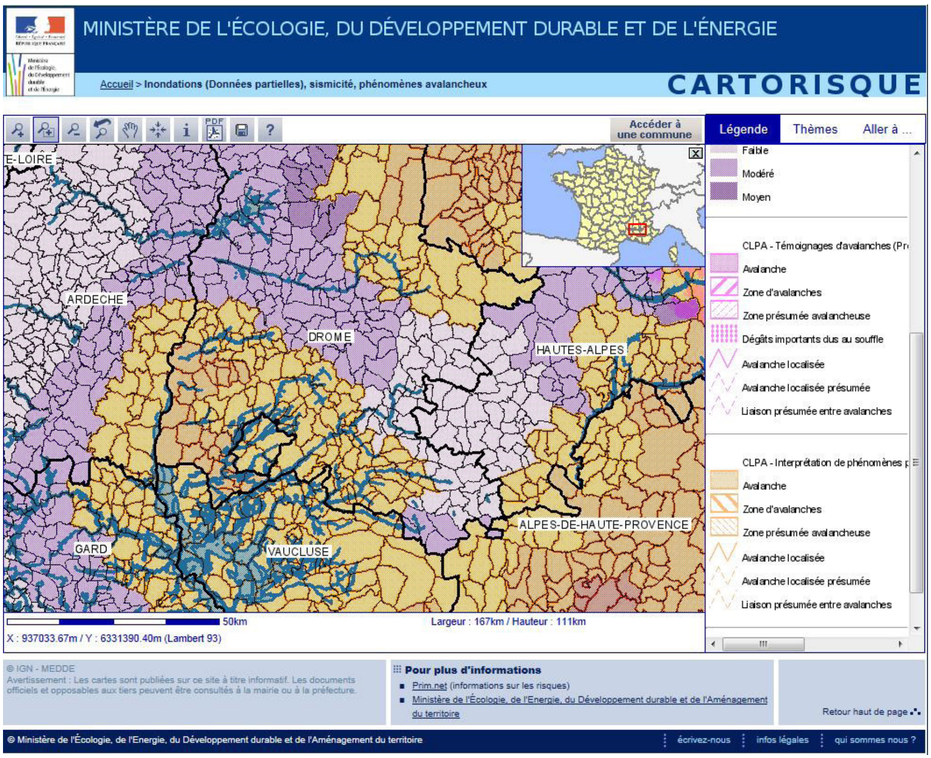

Figure 5: Cartorisque map (Source: Ministére de L'écologie)

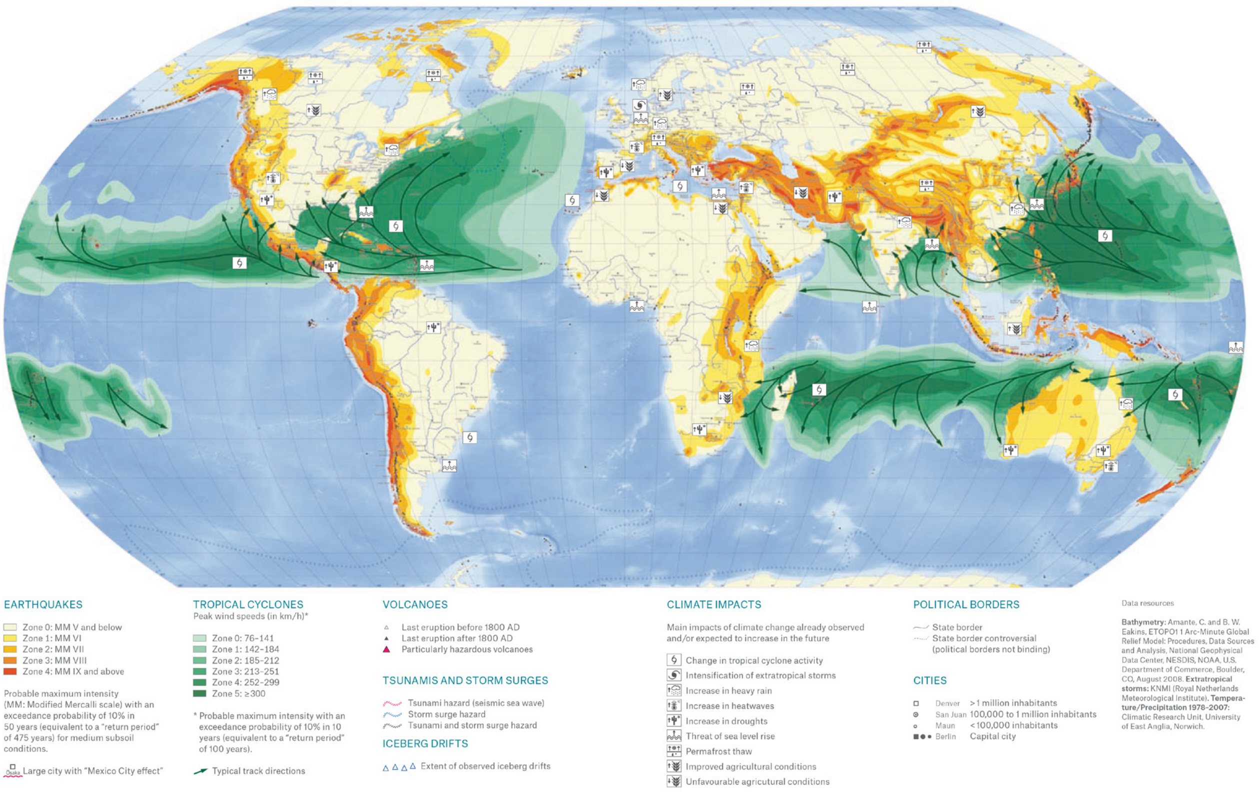

Figure 6: Nathan map of natural hazards (Source: Munich Re)

Another source of hazard information for the perils/regions not covered by vendor models is scientific initiatives, which are typically sponsored by local authorities, scientific and academic collaborations or brokers. This research can have the benefit of identifying risks for events which have not occurred in recent history.

Examples include the Global Earthquake Model, which aims to identify and characterise earthquake risks in regions not covered by vendor models, and scientific publications on flood risk in the Netherlands with the impact of local defence protections.

3.5 Constant vigilance

As noted in Section 2.1, the very nature of NMR presents a challenge to the structured process of identification. Corporate processes are often good at coping with known risks, and known shortcomings, but they may not be as effective when it comes to proactively spotting new sources of risk in time to do something about them.

The authors suggest that there should be a specific function and responsibility for open thinking about NMR. For example, a group or committee could meet every few months to consider what catastrophe risks not already under consideration might be faced by the company, now or in the future. The group would need to be well-informed about the company's plans (see Section 3.2.3) and should include representation from the capital modelling team.

The following techniques may help such a group to continue to stay sharp and regularly challenge assumptions.

3.5.1 Heuristics

Participants will draw from their own experience and use generalisations to help produce estimates when identifying NMR. There should be a common awareness of the flaws in human nature when making such estimates as these are subconscious rules of thumb used to make decisions quickly.

Some common heuristics are:

- **'Anchoring'** to the status quo. It is a natural tendency, for example, to take an offered number, regardless of its relevance, as the starting point for price negotiation. In considering NMR, companies should take care not to start with an expected outcome or a known risk profile. This could occur when someone at the first meeting states that the specific loss from peril X was Y million to the company. It is then easy, but perhaps not relevant, to think about materiality by adjusting from Y.
- **'Availability'** – the tendency to make a judgement about the probability of an occurrence based on how easy it is to recall similar instances. For example, events that have happened recently are given more weight than equally probable events that occurred a long time ago. During the creation of this document, the Thailand floods were probably quoted as an example more often than all other events put together. Extensive media coverage of an event may result in inaccurate perceptions of the contributions of sub-perils, for example, overestimating the contribution to overall loss of tsunami-following-earthquake. In relation to NMR, care should be taken to make sure that thinking is not biased towards events because they have happened recently.
- '**Framing'** – the biased survey. People generally like to agree with statements more than they like to disagree. For example, "Do you agree that the Government is doing a good job?" gets a different response to "Is the Government doing a bad job?" In relation to NMR, care should be taken to ensure that biased questions are not being posed.
- **'Over-confidence'** – the views of an expert in one field of knowledge may be accorded undue influence in considering unrelated fields.
- **'Representativeness'** – the tendency to look for things that are similar and then assume they are the same. In relation to NMR, companies should be sure not to base the probability of an event on what appear to be 'similar' events. They may not be.

4. Framework for quantification

4.1 Introduction

Chapter 1 outlined four categories of NMR. Chapters 2 and 3 suggested governance structures for managing NMR, including ways of identifying them.

Having identified NMR and arrived at an initial estimate of their materiality, the next step is to develop systems for quantifying them. There are many different methods of quantifying NMR, and there should be a structured process for choosing an appropriate one.

This chapter introduces a structure and guiding framework for quantifying NMR, including key considerations when deciding how to measure each identified risk. The methods vary in their complexity and the resources needed for their implementation.

4.2 Quantification methods

For the purpose of Chapters 4 to 6, the authors have categorised methods of quantifying NMR as follows:

- expert judgement – rule-of-thumb methods, often used where little or no data is available, or only a subset of the data is used
- actuarial/statistical – data-driven statistical methods, particularly useful when some loss data is available
- geospatial – methods of varying complexity based on mapping risks within defined geographic areas
- catastrophe model modification and building – using another model as the basis for quantifying NMR

Figure 7 illustrates the four methods in terms of their complexity and resource requirement:

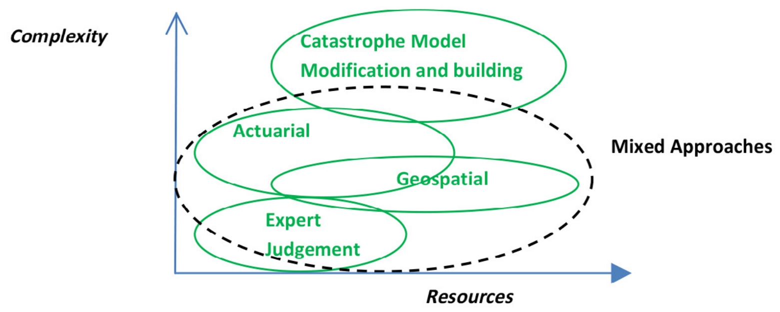

Figure 7: Complexity and resources needed to implement different methods (the intersections between the ellipses are purely illustrative)

This framework illustrates the range of options for dealing with most common types of NMR, and their respective resource requirements (X-axis) and complexity (Y-axis). The ellipses indicate that a given method can have different resource requirements/ complexities depending on how it is used. For example, some geospatial methods such as hazard layers require more resources to create than concentric ring footprints.

The first three methods are discussed in detail in Chapter 5, and the mixed/ modification approaches in Chapter 6.

4.3 Which method to use?

When deciding how to quantify NMR, companies should always take into account the following three considerations:

- what is the purpose of the quantification?
- how material is the NMR?
- what data is available?

There is an additional factor, of course: the resources available. Part of the governance structure for NMR should include communication (as discussed in Section 2.4). If there is a potential mismatch between the materiality of an identified NMR and the resources available to quantify it, this should be communicated to senior management immediately.

4.3.1 Purpose of the quantification

As discussed in Section 2.3.1, while the purpose of the quantification may seem obvious - wanting to understand NMR - the exact nature of the questions to be answered will determine the output required and, therefore, the method selected. Examples may include:

- a company is looking to provide a complete assessment of exposure for the purposes of capital modelling
- investors require clarification on NMR within a portfolio
- a broker is modelling a new portfolio for global access to the insurance markets
- a company is completing its view of risk by considering specific secondary perils
- a government is contemplating an insurance pool scheme
- a company is developing a sophisticated new product

4.3.2 Materiality of risk

The methods used to define materiality (discussed in Section 3.1.1) need not be the same as those used to quantify the NMR. The chosen approach and its required level of complexity should be aligned with the overall exposure to that risk and its materiality within a company's risk appetite. In order to assess materiality, some form of quality data is required; be it premium, exposure or loss.

This paper has already noted that it is very difficult to assess risk materiality without exposure data, or at least a proxy such as market share or premium. An iterative, top-down procedure can be employed initially, whether the quantification method is going to be empirical, geospatial or statistical. This process of refinement will result in conservative exposure estimates in the early iterations but, as complexity increases, more realistic measurements can be achieved. By considering materiality alongside this procedure, the company's risk appetite can be used to define the complexity of the approach.

4.3.3 Availability of data

When determining an appropriate and viable approach towards quantifying NMR, the most significant constraint is generally the availability of data. Data can range from high-level information on premium to detailed exposure data by territory, peril, line of business and/or coverage.

The level of detail required will depend on the type of NMR being analysed. Some perils have a much higher sensitivity to exposure data resolution than others. For example, tsunami risk quantification must take account of large differences in intensity over short distances. This ideally requires building-level coordinate information for accurate modelling. By contrast, in the case of windstorm modelling, aggregates by postcode or CRESTA zone may suffice as the hazard intensity may not vary greatly over short distances.

The type of business being underwritten will also have an impact on the availability of data for NMR assessment. Some books of business lend themselves more naturally to the provision of more complete exposure data - full-value binding authorities, for example - than others such as inwards excess-of-loss reinsurance portfolios.

The availability or otherwise of loss data can also influence which methods may be employed.

IMPORTANT NOTE: A systematic call for data across the organisation often reveals more sources than initially apparent!

4.4 Proportionality

The purpose of considering different quantification methods is to answer the question:

"What is the best way to represent a company's catastrophe risk, taking into account materiality, data availability and the purpose of their investigation?"

The answer may well vary from one risk to another. Therefore, companies should take care to ensure they avoid pre-judging techniques by using emotive descriptions. Terms such as 'simple', 'crude' or 'basic' may be construed as meaning 'bad'. Equally, words like 'sophisticated' or 'complex' can be taken as meaning 'good'. These labels are not helpful.

In reality, proportionality is crucial. The appropriate tool in one context may be quite wrong in another, regardless of how 'simple' or 'complex' it is.

Similarly, there is no necessary link between the complexity of the identification process and the quantification method finally selected. For example, detailed and lengthy investigation may reveal that – taking into account materiality, data availability and purpose – the best way to quantify catastrophe risk in a given region/peril is to take an x% damage factor of aggregate exposures by CRESTA zone. A more complex tool may not add value in this case. A damage factor approach would therefore be appropriate - 'simple' without being 'simplistic'.

Conversely, complexity of calculation does not necessarily confer accuracy of answer.

5. Quantification: standalone methods

5.1 Introduction

This Chapter deals with the first three of the four methods outlined in Chapter 4: expert judgement, actuarial/statistical and geospatial. The latter is split into three, separate sections. The main characteristics of each are described, including circumstances under which they may, or may not be useful. Where appropriate, typical inputs and outputs are listed, examples given, and limitations to keep in mind are suggested.

5.2 Expert judgement

Historically, expert judgement methods were often developed based on experience of a line of business and/or peril in a specific region, on claims data, or by extrapolation of statistical loss studies to regions where there was no data.

In fact, before catastrophe models became ubiquitous, many standard reinsurance pricing tools were developed based on loss information, often expressed as a percentage of a relevant premium or exposure measure. Overall, these are not as complex or resource-heavy as other methods described in this chapter.

Methods that may employ expert judgement include:

- damage factors – applying a damage factor to some exposure or proxy exposure to calculate a loss. This method generally associates the damage factor with a given return period, and a specific distribution in order to extrapolate to higher loss levels.
- premium-to-loss – models where more than one parameter is used to convert a premium-based exposure into a loss-cost. Typical parameters would be a percentage of the premium that produces a reference loss, a base return period, and modifiers to the parameters used for extrapolating to account for portfolio concentration; for example, regional characteristics of the hazard.
- uplift factors – applying an uplift factor to a modelled loss to take into account NMR. This method could, for example, be used to add a given percentage to an earthquake loss in an attempt to account for a tsunami component.

5.2.1 Inputs

The majority of inputs involve the use of aggregate or proxy information to produce an estimated loss. For example, if the total insured value for a particular line of business isn't known, then premium information could be used. For methods involving inflating modelled loss, a modelled-loss-distribution curve is needed as an input.

5.2.2 Outputs

The outputs can include:

- a single modelled loss figure at a given return period in the case of using damage factors against proxy exposure
- a simple approximation of a loss distribution on the basis of a few of such data points
- a modelled-loss-distribution curve which has been adjusted from another distribution

5.2.3 Examples

EXAMPLE 1 - SIMPLE REGION/PERIL MODELS

The single reference loss type model was generally used for perils and regions with one single historical event that was significantly worse than anything else, often alerting the insurance world to the danger from that peril. Expert judgement was used to determine the return period of the event, often based on a study of historical records beyond the insurance industry loss experience; for example, flood heights recorded at key locations. The loss itself was used to determine the proportion of premium/exposure lost in the event, and then a simple distribution – for example, a Pareto - was used to extrapolate the loss to other return periods.

This approach is still potentially applicable to regions where there is very little information and a single material loss has occurred. For example, Cyclone Gonu in Oman in 2007 drew attention to the potential risk from that peril in that region.

Building or licensing a full cyclone model may not be worthwhile. In such situations, an expert judgement approach could be used to assign a return period - for example, to the Gonu loss ratio - and parameters for the extrapolation of the loss to higher return periods could be derived through studies of catastrophe models for other cyclone-affected regions.

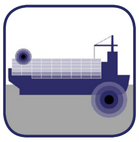

Another example is Typhoon Maemi in South Korea in 2008. Its landfall caused extreme losses to property and casualty classes, as well as marine losses that were largely non-modelled at the time of the event.

EXAMPLE 2 - INTERMEDIATE REGION/PERIL MODELS

Multiple-parameter empirical models attempt to take limited single loss-based models from regions where conditions are different, and apply them to similar regions. The principle behind the method is that there is often independent data about natural hazards that make it possible to quantify the relative frequency and severity in different regions.

Without going into the level of detail required by the third method (see Section 5.4), it is reasonable to assume for some perils that vulnerability to those hazards will be similar between regions. This allows a model to be developed that takes into account known loss experience and hazard variability.

This approach could be used to produce loss curves for hail which take into account the known variation in hail storm climatology between mountainous, flat inland, and coastal regions, based on a couple of reference losses from the market being analysed. Each region would have different parameters, but the same base model approach.

These outputs are often a simple severity measure, such as loss to a portfolio, in addition to a view of the frequency of that loss. To convert this base input into something useful for modelling of the loss process, it needs to be augmented by a method for extrapolating the loss to other return periods. This extrapolation method is often, but not always, an additional component of the expert judgement based loss model.

5.2.4 Limitations

It is sometimes difficult to assess the validity of expert judgement methods. In the case of long-established methodologies, for example, the underlying data and actual analyses that underlie them may no longer be readily accessible.

Solvency II has emphasised the need for transparent, documented methods of validating a company's view of catastrophe risk. Reliance on expert judgement alone has the potential to be seen as a fall-back approach, to be used where other methods are not possible due to lack of data, or are inappropriate on grounds of materiality.

However, the fact that expert judgement has been used – on its own or as a component of another approach – does not necessarily denote weakness or lack of rigour. Nevertheless, there should always be appropriate definition, and documentation, of both the methodology and its application.

5.3 Actuarial/Statistical methods

This Section refers to data-driven statistical methods, specifically where some loss data is available. These methods rely on mathematical concepts, particularly probability, statistics and actuarial science. They may not be appropriate where no loss data exists for the given NMR.

The fact that actuarial/statistical methods are described before geospatial and combined methods does not mean that they are inferior to them. In fact, actuarial/statistical methods of quantifying catastrophe risk in insurance and reinsurance may have been under-utilised as catastrophe models have become more available. In reality, they are an important means of validating all forms of modelling. No model that is inconsistent with actual loss experience should be used unless there is a clear explanation for why the model deviates from historical claims.

There is extensive literature available on actuarial/statistical loss models, and it is out of scope for this document to recommend appropriate methods for all circumstances. Instead, this paper gives a general outline of the data required to perform these types of analyses, and addresses the relevant strengths and limitations of the methods.

5.3.1 Inputs

As inputs, these methods require a statistically useful sample of loss data. Examples include:

- market losses for a given time interval
- company losses for a given time interval. These need to be complete, usually above a given loss threshold. In addition, there needs to be a minimum number of events for statistical fits to be well defined. A reliable record showing very few losses for a long time is still informative. For example, if you know that no earthquakes have occurred in the last 100 years, this is a useful indicator of the 10 and 25-year earthquake loss, but not of the 100-year one.
- data to support adjusting those losses to estimate what the events would have cost for the current portfolio - usually referred to as an 'as-if' revaluation exercise. Typical data used to adjust losses would be current exposure and exposure at the time of the loss, possibly at some geographic resolution. The most appropriate exposure measures would depend on the peril in question, and expert actuarial input is advised at this stage.

5.3.2 Outputs

These methods deliver a statistical model that describes how often events occur (frequency distribution), and also how large the losses are when they occur (severity distribution). These two distributions can be used to generate various types of loss exceedance-probability curves.

5.3.3 Examples

EXAMPLE 1 - STATISTICAL MODELS BASED ON QUALITY DATA

For companies with years of reliable loss-data for a modelled peril such as wind, and a good overview of how their portfolio has changed, it is worth building a statistical loss model to compare with the low return period results from the catastrophe models.

The process involves:

- Updating the raw historical losses for exposure growth, inflation, building codes, insurance take-up and changes in insurance coverage – the 'as-if' correction
- Determining the threshold above which the loss record is complete. This threshold is often a far higher value than the actual loss value that was chosen as the threshold for recording event losses
- Taking the known time period and the list of 'as-if' losses, it is possible to fit a statistical model that describes how often losses occur (frequency) and how large they are when they occur (severity). These two elements can be used to generate a full curve of aggregate or event losses along with their exceedance-probabilities. In addition, there are a number of statistical techniques - for example, bootstrapping- that can be used to estimate confidence intervals around the statistically modelled curves.

EXAMPLE 2 - STATISTICAL MODEL BASED ON INDUSTRY-WIDE INFORMATION

For certain regions/perils, industry databases have been constructed to collate exposure and loss information. An example is the Insurance Council of Australia (ICA). Some of those perils/regions may not currently be included or licensed in the available catastrophe models. The industry loss-data for the non-modelled perils can be corrected to account for growth in exposure, and a statistical model fitted to them. This fitted reference curve, for example, based on Australian bushfire losses - can then be adjusted according to the company's market share to represent the non-modelled peril.

5.3.4 Limitations

These methods have many limitations, which, historically, have driven the development of catastrophe models. The key drawback is that reliable loss records are often too short to be of practical use. There are many possible reasons for this, ranging from simple lack of recording of losses to very complex changes in portfolios and policy conditions that render recent event losses difficult to trend.

Also, the inherent assumptions in the statistical methods need to be understood in terms of the loss processes being modelled. Most statistical methods assume that the loss events are independent of one another, and derived from the same statistical distribution. In practice, there are many ways in which catastrophe losses can violate these assumptions; for example, if there are loss inflation drivers that affect large losses more than small ones, extrapolation from the smaller losses will underestimate the potential extreme events.

Other examples of limitations include:

- clustered loss processes where frequency and severity are linked
- cyclical loss processes where the observed period may have anomalously high or low loss activity
- poor sampling of the range of possible intersections of hazard and exposure, either due to high concentrations of exposure, or significant new exposures that have appeared since the beginning of the loss observation period
- grouping of fundamentally different loss processes together. For example, if losses from local convective summer storms and large scale winter storms are combined, the statistical assumptions will be violated. This is an important point, as there is often a temptation to include all significant losses from a single broad hazard category to increase the amount of available data.

5.4 Geospatial methods

Geospatial methods involve a group of processes of varying complexity that have been developed either within organisations themselves, or which already form established best practice of professional organisations such as Geographic Information System (GIS) associations and practice groups.

Advances in IT and related mapping techniques over the last twenty years have encouraged increasingly sophisticated geospatial methods to evolve. One thing they all have in common is their reliance on GIS technology and defined geographic boundaries, such as country and province borders, CRESTA zones, buffers, etc.

This next Section discusses three approaches. In order of increasing complexity they are:

- Geospatial methods that only utilise geographical boundaries
- Geospatial methods that utilise custom shapes
- Geospatial methods that utilise hazard layers

The flow diagram below in Figure 8 illustrates the thought process associated with geospatial methods.

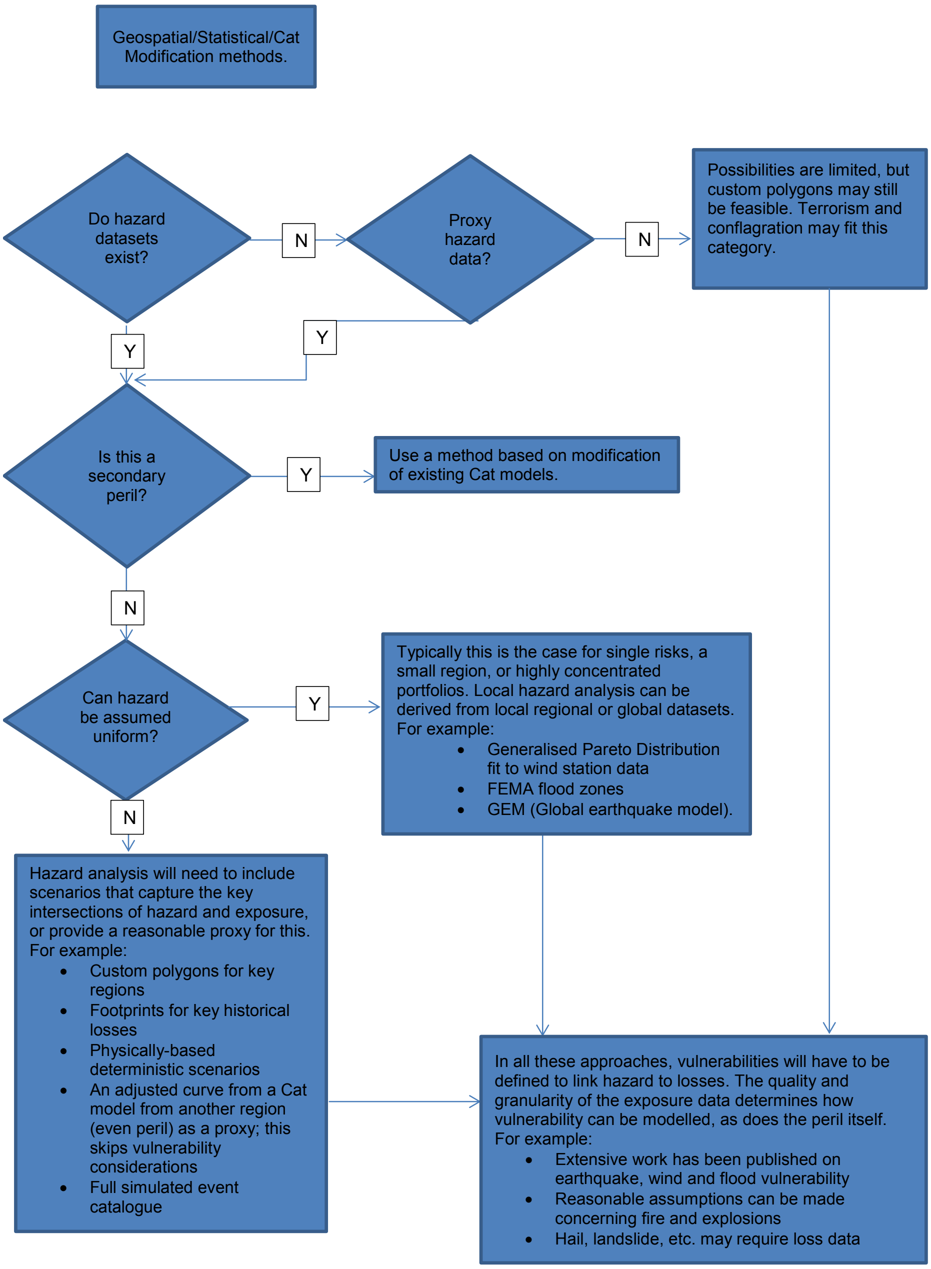

Figure 8: Geospatial flow chart

5.5 Geospatial methods that only utilise geographical boundaries

There are several relatively well-established geospatial methods that only utilise geographical boundaries. They largely trace their roots to the original PML techniques, where exposure in a geographical area was accumulated either in totality or by line of business, and then multiplied by a percentage to produce the PML.

Historically, these percentages (sometimes confusingly referred to as PMLs themselves) were derived from underwriters' experience and local practices, therefore they have much in common with the expert judgement methods discussed in Section 5.2.

Over time, these methods have evolved through comparisons with output from catastrophe models. They are now an integral part of the Standard Formula approach adopted by Solvency II where Q-factors (Solvency II factors indicating implied 200-year damage ratios) are applied to different countries and perils to represent the 200-year loss (i.e. the loss exceeded on average every 200 years or more) for that country and peril.

The Q-Factor approach can be extended to smaller geographical units within countries such as provinces or CRESTA zones. The losses at a given return period can be combined in a similar manner to the square root formula given in the Standard Formula. For example, a country-level 200-year loss could be the square root of the squared sum of the product of the factors and the corresponding total insured value.

5.5.1 Inputs

In the example below, inputs would be total exposure/premium by CRESTA and line of business (LOB), and a damage ratio also by CRESTA/LOB. Geospatial damage ratios can be assessed through asking underwriters or via benchmarking to a specific low threshold at a given return period.

5.5.2 Outputs

A "loss" estimate at a defined return period by CRESTA/LOB.

5.5.3 Examples

EXAMPLE 1 – ISRAEL EARTHQUAKE

This map indicates possible 200yr damage ratios by CRESTA in Israel.

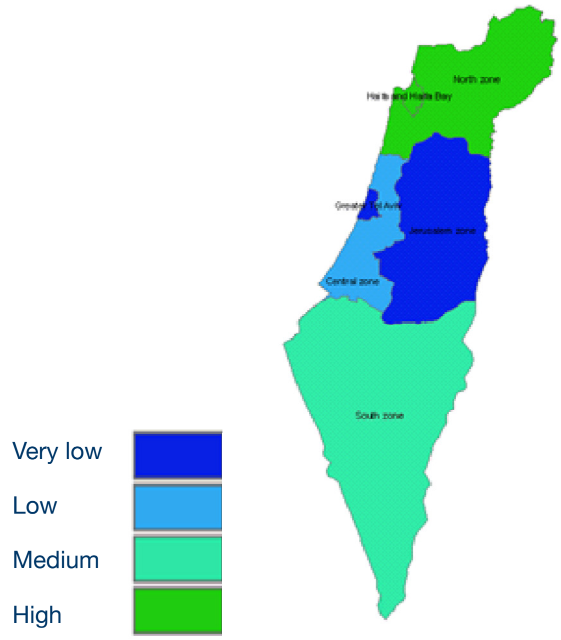

Figure 9: Israel earthquake map

5.5.4 Limitations

When considering damage ratios, the size of the country and proximity of exposure to risk should be taken into account. For example, two countries of similar earthquake risk could be assumed to have similar building vulnerability, but one is significantly smaller than the other. Whilst both may have similar seismicity, the smaller country can be expected to have a higher damage ratio applied. Its smaller geographic extent creates a lower likelihood of exposure being away from earthquake risk compared to the larger country.

5.6 Geospatial methods that utilise custom shapes

Next up in complexity are methods which rely on abstract shapes that may have a relationship to actual geographical boundaries, but need not. A common term for these shapes is Custom Polygons (although from a pure geometric point of view they do not have to be polygons), which will be used for the remainder of this paper.

These geospatial methods typically work by creating custom polygons – meaning shapes such as circles, ellipses, n-sided polygons or coastal boundaries – and overlaying them on exposure data. There is no reference to hazard information.

Below are examples of three potential NMR and possible approaches to their quantification using this technique. Please note that the shapes described are only examples, and that real situations and application could be quite different.

EXAMPLE 1 - TERRORISM BLASTS, METEORITE STRIKES OR VOLCANIC IMPACTS

Accumulation techniques involve placing circles of varying radii around a location, with corresponding damage ratios, as in Figure 10, below. The damage ratios will tend to be higher for those circles closer to the blast/strike/impact than circles further away. Estimated losses can then be passed through a company's own financial module or vendor catastrophe model.

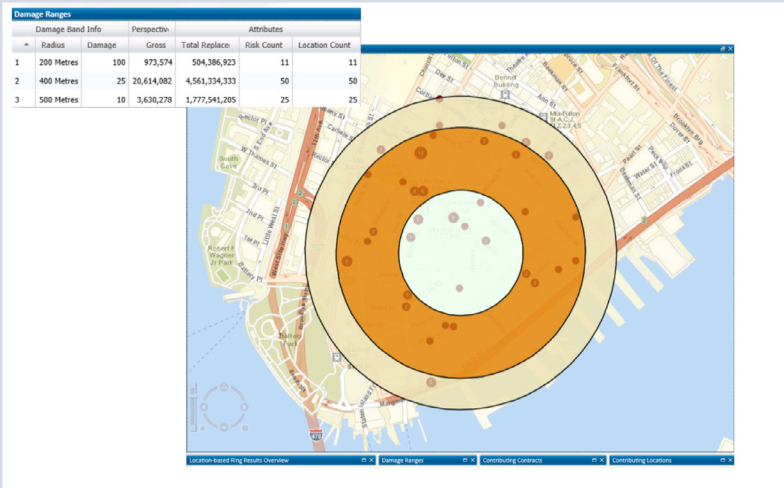

Figure 10: An example of quantifying terrorism risk using concentric circles with different damage ratios by distance (Source: AIR)

Once again, materiality of exposure should determine the focus of investigation. For terrorism, analysing key cities or landmarks by replacement value will identify key areas of potential loss. For volcano risk, which is highly location-dependent, checking high materiality to the portfolio for Europe could start with areas around Mt Vesuvius in Naples and Mt Etna in Sicily.

A relative terrorism risk heat map is shown, in Figure 11, illustrating largest exposure concentrations and 400-metre accumulations of risk. The inputs for this method could include latitude, longitude, replacement values, and policy terms for all locations in a region. On the hazard input, these could be a GIS shapefile. A geospatially varying damage ratio would also be applied to the exposure.

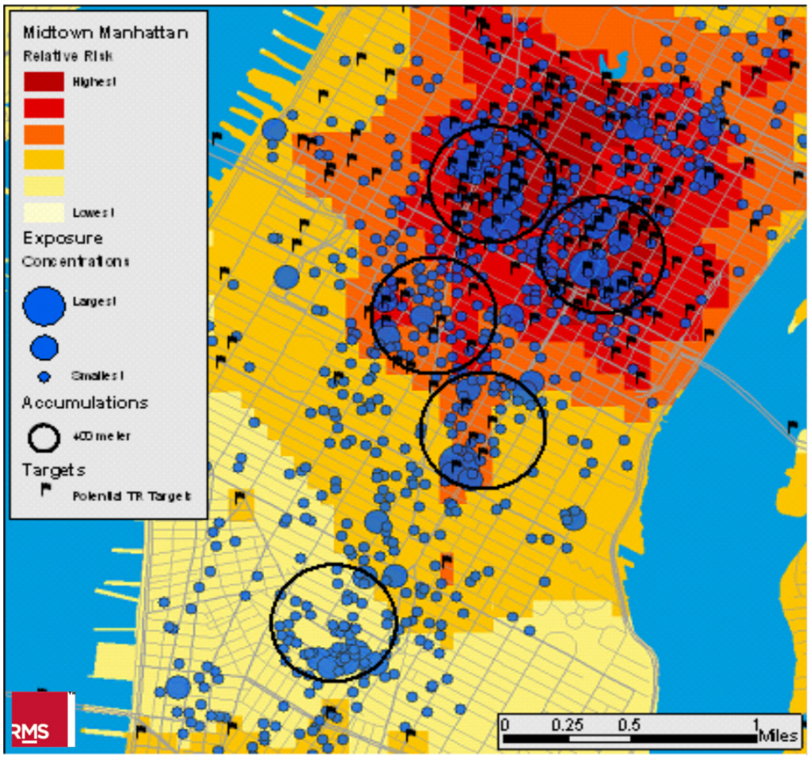

Figure 11: Terrorism heat map (relative risk heat map for terrorism risk in Midtown Manhattan - this map includes an overall view of the largest exposure) (Source: RMS)

EXAMPLE 2 – TSUNAMI, STORM-SURGE

An approximate calculation of tsunami risk exposure can be accomplished by creating a tsunami footprint as a custom polygon that runs parallel to the coast - for example, 3km from the coast, assuming inundation up to that distance - which overlays the underlying exposure. The damage ratio can be banded such that higher damage ratios are applied to locations closer to the coast and lower damage ratios further inland.

EXAMPLE 3 – TROPICAL CYCLONE

Tropical cyclone footprints can be approximated by bands of parallel lines. For example, the inner bands would have higher damage ratios than the outer bands. The above shapes can be moved geographically to give a range of possible losses by changing the geographic location but keeping the damage ratios the same.

5.7 Geospatial methods that utilise hazard layers

The next group of geospatial methods takes into account the hazard information. Typically, this requires creation of GIS layers relating to specific hazards, and is, therefore, more complex than the custom polygon method.

The table below gives further examples of four potential NMRs and possible approaches to their quantification. Again, the shapes listed are only examples, and real situations and applications could be quite different.

EXAMPLE 1 – TSUNAMI

A tsunami footprint can be developed using a wave propagation model and the associated bathymetry and topography for the region under consideration. This can be seen as an extension of Example 2 in the previous section, which had a tsunami footprint developed largely on the basis of expert judgement (for example, 3km from the coast).

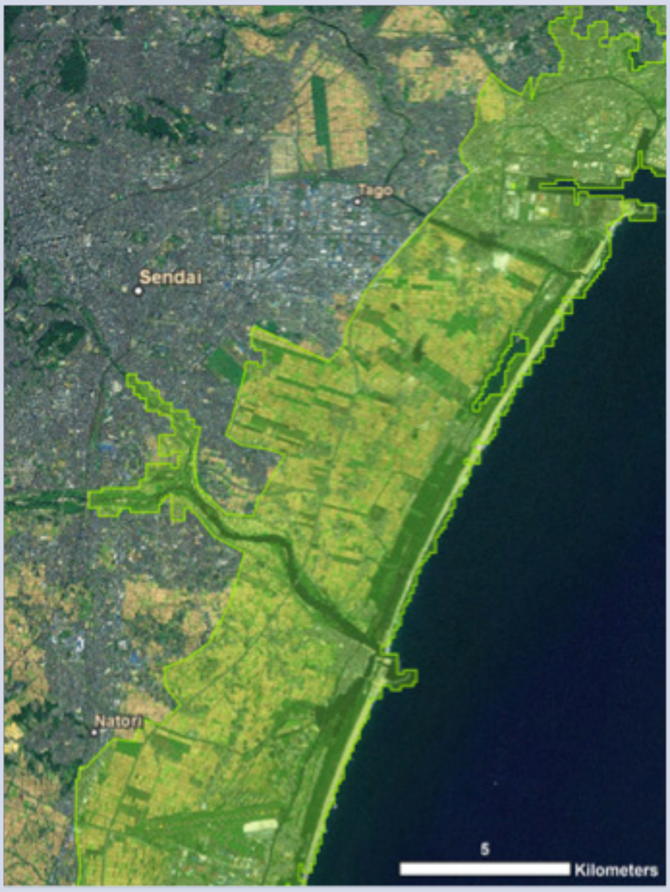

Figure 12: Tsunami footprint (Source: AIR)

EXAMPLE 2 - FLOOD ZONES

These are designated regions based on the underlying susceptibility to flooding. Flood zones can be sourced from various national reference authorities; for example, the Federal Emergency Management Agency (FEMA) in the USA. They could be based on either hydrodynamic simulations and/or historic observations. A FEMA flood map can be overlaid on exposure data to understand the proportion of exposure susceptible to different degrees of flooding.

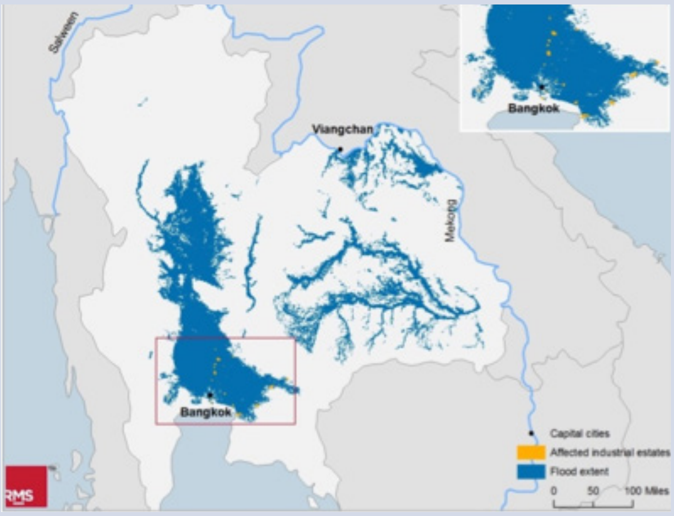

Figure 13: Thai flood map (Source: RMS)

EXAMPLE 3 - LIQUEFACTION

Liquefaction maps can be used to accumulate exposure which has susceptibility to liquefaction. This can then be used to potentially inflate modelled earthquake losses in regions where earthquake damage is modelled without liquefaction.

EXAMPLE 4 - EVENT FOOTPRINT

The wind footprint of a historical storm such as Hurricane Katrina could be used and overlaid over exposure to determine the total exposure at risk. This footprint can be moved geographically to understand the potential range of losses, for example, by changing the landfall location. A similar exercise could be done for earthquake risk using ground-motion footprints.

5.8 Combined methods

Any or all of the techniques discussed above can be used in combination. For example, modelling terrorism blast may involve geospatial methods to create a defined footprint such as a circle, plus damage ratios derived from expert judgement.

Once again, proportionality is crucial. ‘Complex’ techniques can be used alongside ‘simple’ ones in quantifying a single NMR, according to the materiality and data-availability for different risk components.

A probabilistic catastrophe model could be considered the most complex example of a combined approach. However, the techniques involved in building a region/peril catastrophe model fall outside the scope of this document.

6. Quantification: extending cat model use

6.1 Introduction

This technique applies when a company already has a model for a catastrophe risk that may be similar to the identified NMR. In this case, modifying the inputs or outputs – or both – of the existing model may give an appropriate representation of the NMR. This has the additional attraction of leveraging the company's existing capability and licenses.

However, the method is only possible when the NMR has some relationship to a catastrophe risk that a company already quantifies. For example, the effects of a secondary peril such as tsunami may not be represented in a company's existing earthquake catastrophe model, so the outputs of that model could perhaps be modified to take account of the secondary peril. Similarly, the loss curve for windstorm in one country could perhaps be used as the basis for quantifying non-modelled windstorm risk in another country.

The existing model will often be a probabilistic catastrophe model, but need not be.

Examples of modification approaches include:

- modifying the exposure-value inputs in an existing model to account for NMR
- adjusting existing model output

6.2 Exposure modification

This approach involves adjusting the exposure values that are inputted to an existing model to take account of risk that is not represented. There may be a number of reasons why the exposure-data is under-stated, including the following.

6.2.1 Non-capture

There may be exposure-data that has insufficient geographical identification to be modelled, for example, where a non-existent postcode has been used. Rather than omit these values altogether, geo-coded exposure data can be inflated by the proportion of unidentified exposure. This ensures the total insured value is modelled, though uncertainty in losses would be brought about by some exposure not being imported to the correct location, wherever it may be.

Similarly, potential under-insurance can be accounted for by grossing-up exposure values.

6.2.2 Non-modelled classes/lines of business

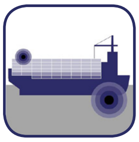

For some mobile coverages such as cargo and specie, assumptions could be made that a percentage of the exposure is within the footprint of an event.

Other approaches could involve using a "worst case scenario" where non-modelled lines of business such as marine/pleasure boats are assumed to be on the coast in the path of a storm event.

6.2.3 Coverages

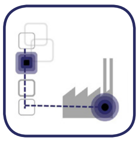

Supply chain risk can be quantified by finding the total dependency to suppliers - for example, first and second tier - and approximating the exposure to such risk under failure of these suppliers.

Such risk can also be accounted for by incorporating this inflated business interruption exposure into the model.

6.3 Adjustment to existing model output

Adjusting the output from an existing model may be suitable for:

- capturing non-modelled secondary perils where a model exists for the primary peril
- capturing non-modelled sources of loss such as loss adjustment expenses
- making adjustments to the primary peril event losses to capture non-modelled aspects of the loss process, for example, temporal clustering

6.3.1 Adjustments to probabilistic models

Although the outputs from catastrophe models are commonly thought of as curves, meaning exceedance-probability curves, they actually consist of data-sets called event-loss tables or year-loss tables.

Detailed discussion of these tables is outside the scope of this document. It is sufficient to understand that adjustments to the loss-values or loss-probabilities in these tables will affect the exceedance-probability curve. Loads can be applied by adjusting the frequency, the severity or the variance of events in the output, or a combination of all three.

- frequency adjustment can be used to incorporate temporal clustering (see Section 3.3.1), or to emphasise or de-emphasise events with a specific characteristic. Frequency adjustments are generally the most straightforward to apply. Depending on the type of vendor model, output adjustments can be applied to the event frequency of an event, or by resampling a simulation. Resampling a simulation to alter the frequency requires removal or repetition of existing events, and may remove any clustering built into the vendor's outputs. These adjustments can be applied within the company's capital model or at a prior stage.
- non-modelled secondary perils can be accounted for by severity adjustments, which are more complex. The purposes of a severity adjustment include increasing an event's mean damage ratio, accounting for a larger exposure footprint than the existing event, and adjusting for non-modelled sources of loss. The approach taken will depend on the nature of the NMR. Particular care must be taken not to adjust the ground-up loss to a level beyond any contractual limit, particularly in the tail of the distribution as this will result in potentially material overstatement of losses.
- variance adjustments can be applied where secondary uncertainty has been preserved in the catastrophe model output. The absolute size of the standard deviation can be altered, or the mathematical distribution can be changed within the capital model. Generally, any severity adjustment should take into account the variance and the decision of whether to scale it relative to severity. Following adjustment of the severity to account for a non-modelled risk, this may increase or decrease uncertainty around the mean severity.

6.3.2 Parameterisation

A combination of expert judgement, actuarial/statistical and geospatial methods described throughout the previous Chapter are all applicable to parameterising a secondary NMR. Secondary risks have additional complexity due to dependence on the underlying hazard and the potential to double-count the loss from primary/secondary perils at a site.

6.3.3 Dependence on primary peril

A proportionate or fixed adjustment across all events in a probabilistic loss output may not be representative of the NMR interaction with the primary peril. For example, liquefaction hazard following earthquake is conditional on the presence of liquefiable soil, depth to ground water, and a threshold severity of ground shaking and duration. In such cases, the adjustment should be made to a subset of events. Choosing the subset can be informed from empirical or geospatial sources, and information on the event location and magnitude from the vendor tables.

Another example is 'demand surge'; the short-term increase in labour and material costs after a catastrophe event. One means of quantifying the potential short-term inflation in losses is adjusting the outputs of a probabilistic model beyond a particular return period. The exact return period could be decided as having an insurable/economic loss greater than a particular percentage of GDP for that country, for example. This would allow the chosen return period at which losses are amplified to be variable depending on the stress of the primary peril to the country. Another calibration-point could be historical experience. For example, if demand surge has been observed in a country once in the last 20 years, should it impact losses at a return period much more frequent than that?

6.3.4 Double counting

If a non-modelled secondary peril has losses that are correlated with the primary peril, care should be taken to avoid double counting. Validation should be performed on the combined effect of all sources of loss from an event rather than each peril in isolation. An example of this is hurricane winds and ensuing storm surge. Since the hazard is highest at the coast for both the primary and secondary peril, loading a wind-only model to include the effect of storm surge needs to be loaded by the marginal impact of the surge, rather than the standalone effect.

6.3.5 Effect of policy conditions and limits

With severity/exposure loads, the point in the model where the load is applied will affect the final impact. In most of the hazard examples discussed in this document, adjustment should be performed on ground-up loss, so that the effect on gross/net losses can be ascertained. However, this will not always be the case, and loss adjustment expenses may apply in addition to any policy exposure.

6.4 Example

ADJUSTING AN EXISTING OUTPUT – DAM FAILURE

A mountainous region with many dams may be prone to avalanche and subsequent dam failure. The cause may be a long-lasting rain event, or an earthquake. Therefore, dam failure may perhaps be appropriately considered as a kind of secondary peril, linked to events in an existing primary probabilistic flood or earthquake model.

The assumption with this approach would be that event footprints in the primary model act as proxies for trigger events and increasing loss outputs for these events may be a valid representation of dam failure. Alternatively, the same events could be used to trigger a custom-built adjustment to the loss outputs tailored to the event itself, as well as the geographical and topographical features of the exposure.

In a very detailed assessment, a modeller might choose to assess specific deterministic events with detailed location-level assessment for the event impact and possibilities of mitigation or early-warning, all of which could also be included in a probabilistic portfolio level approach.

6.5 Summary of quantification outputs

A number of different outputs can be produced from the methods described in Chapters 5 and 6. These include:

- simple losses to portfolios, possibly accompanied by an estimation of the return period or frequency of the loss level that may be extrapolated to other return periods
- severity and frequency distributions to describe losses, as well as loss exceedance-probability curves to describe the probability of losses to portfolios
- geographical exposure totals by geographical region such as CRESTA, or calculated for a custom polygon
- geographical loss totals by region such as CRESTA with an estimation of the return period or frequency of the loss level, or calculated for a custom polygon
- loss outputs in the form of exceedance-probability curves extrapolated from existing loss curves, with a description of the adjustment basis
- total exposures for a particular geography, adjusted, along with a description of the adjustment basis
- full probabilistic output, consisting of complete loss tables

All these outputs can be used as inputs to capital models, as described in more detail in Chapter 7.

7. Incorporating NMR into the Capital Model

7.1 Introduction

As stated in earlier chapters, NMR should already be included implicitly in the capital model. This Chapter discusses ways of explicitly representing these risks, and considers issues that may arise when doing so. Good governance will be discussed, as will correlation between NMR and other components of the capital model. This Chapter finishes with a brief description of how to test the capital model output, including examples.

7.1.1 Definition

Section 1.2 provided a definition of NMR.

As we have seen, inadequate representation may derive from any or all of the four categories of NMR. Assessing 'adequacy' at any point must take account of risk materiality, evolving statistical methods and standards, new science, better information technologies and – perhaps most of all – messages from the real world about the nature, frequency and human consequences of catastrophes.

In this context, NMR does not necessarily imply that a vendor catastrophe model has or has not been used. As discussed earlier, there are circumstances where it would be entirely appropriate not to use one.

7.1.2 Scope

This chapter assumes that:

- the internal economic capital model is based on Monte Carlo simulation, which serves as the risk-consolidation platform for all the risks to which a company may be exposed
- the principal inputs to the capital model for catastrophe and terrorism are so-called event-loss tables or year-loss tables, from which individual loss cash-flows are drawn. These tables are the product of a probabilistic catastrophe model or similar tool that represents a discrete, 'empirical' cumulative loss distribution function specifying the likelihood that a given loss severity is exceeded.

Quantification of dependencies is outside the scope of this document.

7.1.3 Team structure

While there are many possible team structures inside companies, this Chapter refers to the 'catastrophe modelling team' and the 'capital modelling team' as separate entities. In reality, the two functions can be handled by one team or even the same person. However, the distinction is helpful for the purpose of distinguishing between functions in the workflow.

7.2 Governance

The capital model sits at the end of the process of identifying and measuring NMR. At some point, the decision has to be made as to how NMRs are included in the capital model.

As we have seen in Chapters 3 and 4, a clear, structured process is essential for identifying and quantifying NMR. This applies with equal force at the capital model end of the process.

The methods chosen to represent NMR in the capital model can range from a simple distribution to more sophisticated event-loss or year-loss tables. The method chosen will depend on many factors, including the type of risk, its materiality, and who is responsible for making decisions about perils to include and methods to use. The responsibility for making these choices, and the processes for implementing them, must be clearly defined and articulated.

We have already seen that resources can be an issue, and that working on NMR can be time-intensive. This also applies to the capital model. The processes should explicitly contain elements that resist downplaying NMR as negligible without adequate investigation.

7.2.1 Communication

The capital modelling team has its own communication requirements for managing NMR. The following pointers may be helpful:

1. The NMR process should ensure that the capital modelling team is aware and involved early on.

   Nothing will be gained by getting to the end of an expensive and time-consuming quantification process only to discover that the outputs are of no use to the capital modelling team. The question “What sort of data is the capital modelling team going to need?” should be considered throughout, and the easiest way to find out is to ask them. The capital modelling team should be involved in parameterising risks within the capital model.

2. Data given to the capital modelling team must include clear descriptions of what is included and, just as importantly, what is not.

   Passing data between teams can be problematic. The better the communication between business, catastrophe modelling and capital modelling teams, the less risk there will be of data issues. Data should generally be proportionate to the risk and its materiality. In other words, if the data is too detailed, it can lead to false confidence and spurious accuracy.

3. The process should allow for corporate growth and its effects on team structure.

   As noted in Section 7.1.3, a company’s catastrophe modelling and capital modelling functions may be housed in one team. If the company grows, however, the structure may develop into multiple teams. The process should, therefore, explicitly recognise the need for communication across different functions, regardless of who is performing them at present.

4. The process should ensure a close working relationship between underwriting, catastrophe modelling and capital modelling. For example, companies should ensure that the capital modelling team is represented in any forum that specifically considers NMRs.

5. The process should explicitly recognise the need to communicate uncertainty around NMR. Uncertainty is often higher for NMR than for other risks, and this should be communicated to senior management, including the related impact on the uncertainty of the overall capital figure.

6. The process should explicitly recognise wider issues in the company.

7. NMRs do not exist in a vacuum, and recognising them explicitly may have consequences for the company. For example, an increase in the capital requirement may be resisted by underwriting or management, but these issues can be minimised by involving these stakeholders early on.

   The process should specify how results are communicated to senior management and the board. When the results from the capital model are reported, it may be useful to highlight the work that has been done on NMR. If material enough, some of the assumptions made should be explained in a clear, non-technical language.

7.3 Inputs

When referring to a 'full probabilistic curve', this includes either a distribution or an event, or year-loss table.

For NMR to be incorporated, full distribution assumptions are needed for the stochastic model. Typically, this will either be an actual distribution including suitable parameters – for example, a Pareto distribution with shape and scale parameter – or an output table similar to, or derived from, the output of a catastrophe model.

What is provided to the capital modelling team for NMR may fall short of this requirement. For example, there may only be a single return period damage factor. In such cases, the techniques in Chapters 4, 5 and 6 may be used to turn what has been provided into a full distribution. This will probably require working with underwriters, catastrophe modellers and other experts.

However, there may be cases where no full distribution can be produced. What should companies do when the least possible amount of information has been passed on to the capital modelling team (for example, a single scenario with associated loss severity and frequency)? This data point may have to be used to parameterise the NMR.

7.3.1 Single point estimate

To convert the single point into a full probabilistic curve, several approaches are possible, including:

- benchmarking against industry losses may be suitable to generate a curve going through the one point available
- the shape of the distribution could be determined by expert judgement or through comparison to similar perils, and the curve could then be fitted through the single point
- a single-parameter Pareto curve can be used, although this is likely to exaggerate the tail losses

EXAMPLE 1 – EARTHQUAKE

Based on expert judgement (see Section 5.2), one loss and one probability point have been estimated for earthquake in Country A. Using techniques discussed in Chapter 3, the exposure has not been deemed very material, so the company has decided to focus on other perils rather than investigate this specific risk in greater depth.

The capital modelling team has access to a shape of distribution that has been produced from a full probabilistic catastrophe model for the significantly more material exposure in neighbouring Country B. This can be used as an approximation, with the curve being fitted through Country A's single loss/probability point to create a full distribution.

A more detailed quantification of correlation parameters is outside the scope of this Chapter. For techniques to convert a small selection of damage factor scenarios into a full distribution, please refer back to Chapters 5 and 6.

7.4 Issues to consider when including non-modelled perils

This Section will highlight some areas where special attention is needed for NMR in the capital model. Dependency, which is one of the most significant areas, will be discussed in the next Section.

7.4.1 Materiality

While materiality has been extensively discussed in previous chapters, it is important to mention that what is material for the catastrophe exposure of the business may not be as material in the capital model due to diversification.

However, while individual perils might not be significant contributors to the capital requirement, a set of potentially correlated NMR overall could have a material impact.

Therefore, it is essential to investigate dependencies between NMR as well as interaction with modelled areas to fully understand the materiality of a risk to the business.

EXAMPLE 2 – CONTINENTAL FLOOD

A company has estimated its exposure to flood in Eastern and Central Europe. As the exposure is less material than for other perils, damage factor approaches are used to estimate the risk by country. However, correlation between countries – and hence overall risk – has been underestimated, as has the impact on capital.

During a later update of the capital model, correlation across territories has been reassessed. The overall impact on capital has been more material than previously estimated, leading to a rise in capital required.

In both of these situations, the opposite effect could of course also be true.

7.4.2 Adjusted event-loss or year-loss tables

A common approach, discussed in Chapter 6, is loading inputs or outputs from a vendor catastrophe model. We have seen that options for modification include:

- alteration of the exposure before the vendor catastrophe modelling process step
- alteration of any commercially available vendor model, for example, in its vulnerability module
- scaling of the values in the losses - frequency, mean severity, secondary uncertainty distribution

It is essential to communicate which loading has been applied to avoid potential double counting or omission of loading factors.

Most capital models allow the users to increase events 'on the fly' by a certain factor, and this often can be the simplest method.

There are several common mistakes to avoid:

- severity loading factors should in most cases be applied before the application of policy conditions or, if that is not possible, at least before reinsurance structures (inwards or outwards) are applied
- loading should not necessarily be even along the exceedance-probability curve, and while a fixed factor is often the easiest solution, it might not be the most appropriate
- results will differ if a frequency loading is used in comparison to a severity loading. Therefore, it is essential to understand the differences, and to decide which one to use. An important caveat with adjustments to the frequencies of stochastic events is that the resulting event realisations lose their meanings as scenarios.

EXAMPLE 3 - LOADING FOR SURGE

A company writes an inwards reinsurance Excess of Loss treaty covering US hurricane exposure. The contract has a $10m limit applied. As the broker-supplied results do not include storm surge, the company decides to increase the loss output by a fixed factor for all event IDs.

However, as this has been applied after the treaty limit, it leads to single event losses in excess of $10m dollars - outside the maximum possible loss to the treaty - potentially leading to an unrealistic capital requirement.

7.4.3 Model change

It is also essential to consider non-modelled perils in the model change policy, at least implicitly. Improving the quantification of non-modelled areas could potentially lead to change, up or down, to the solvency capital, so it is important to understand potential impacts early on to avoid unexpected and sudden changes.

Good risk quantification should take priority over the potential workload associated with a model re-submission or capital adjustment.

7.4.4 Hours clauses

Depending on the catastrophe model, if any, used, it might not be possible to model hours clauses before the capital modelling stage. This is just as relevant for non-modelled perils, especially as the quantification of loss is often more simplistic, and hence hours clauses will most likely not have been considered.

Depending on the capital model used, it might be possible to model hours clauses, which can cause significant changes for reinsurance contract losses. As a result, it is important to communicate to the capital modelling team if these clauses have been considered, and to estimate if it is material in order to try to implement these inside the capital model.

7.4.5 Worldwide or ’other risks’ loss curve

Chapter 3 mainly focused on identifying individual NMR for the business. For non-material perils, it is common practice to also include a loss curve which tries to encapsulate all missing items. This distribution will often be based on historic losses or be benchmarked against industry losses.

It is good practice, therefore, when improving modelling of specific NMR to ensure that this ‘other risks’ curve gets adjusted accordingly to avoid double counting. In addition, the methods discussed in earlier chapters might be useful to improve the level of sophistication when estimating the overall distribution.

7.4.6 Shape of the distributions used

From a capital point of view, some of the methods used in previous chapters, may not focus on the most important parts of the curve.

Regardless of the method, it is important that the capital modelling team understands the decisions and model points chosen when quantifying NMR to ensure that there is a suitable alignment between what has been delivered and the capital model.

EXAMPLE 4 - DAMAGE FACTOR APPROACH

A straightforward damage factor approach has been used to estimate return periods up to 1-in-100-years for underwriting purposes, and to quantify realistic losses the company might expect in a lifetime.

However, the capital model uses a 1-in-200-year return period. In addition, the company tries to adjust its overall 1-in-100-year Tail Value at Risk (TVaR) loss for group-wide enterprise risk management purposes.

In this case, the damage factor estimations might not be suitable for the tail end of the curve. The methods in Chapters 5 and 6 are equally useful to correct the area of the curve dealing with highly unlikely risks. Alternatively, comparison with modelled perils or expert judgement might be helpful.

7.4.7 Attritional losses

As with Section 7.4.5, a more sophisticated method of modelling NMR might overlap with the actuarial methods used to calculate attritional losses.

The capital modellers should be aware of what has been included so that adjustments can be made either to the non-modelled curves or the attritional losses to avoid double counting.

7.4.8 Business forecasting

The capital model may project exposure forward into the next business year, while the catastrophe modelling group focuses on exposure and loss estimates for the business written today. NMR can increase significantly over the course of the business year if the company has chosen to invest in a certain new territory.

Again, communication is the key to ensure that there is a feedback loop between business forecasting, capital and catastrophe modelling so that perils which may become material to the company can be assessed early on. This avoids sudden surprises through losses in new territories as well as unexpected jumps in the solvency capital figure in the future. As a first order approximation, the planned change in premium income (which will be known in most such situations) can be used as a proxy for the change in exposure and risk from the growing parts of the portfolio.

7.5 Dependencies and correlations

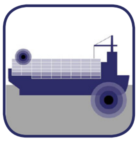

Understanding the relationship between losses to different lines of business and different potential events is essential when dealing with NMR.

We use the word ‘dependency’ to describe this relationship, rather than ‘correlation’, which has a specific actuarial and mathematical meaning. ‘Dependency’ will indicate any relationship between elements or processes that were previously assumed to be independent.

For example, the World Trade Centre loss of 9/11 revealed a dependency between classes of business - Aviation and Property - but only in a catastrophic event. Under normal circumstances, these classes really are fairly independent.

Dependencies will somewhat be determined by how a company defines region/peril. For example, ‘European Flood’ will have different dependencies from separate Germany, Poland and Austria Flood.

7.5.1 Dependencies across lines of business

By definition, catastrophes are rare events, and when they do occur, they may be expected to produce losses from multiple lines of business. The following dependencies should be considered:

- what is the relationship of frequency of losses between lines of business for a given region/peril? Should we assume 100% dependency? If one line of business suffers a loss, do all the others?
- what is the relationship of severity of losses between lines of business for a given region/peril? Do we need to determine the level of dependency between lines of business? Will this vary by region/peril?
- what is the relationship of aggregate losses between region/peril? Is it dependent upon the definition of the ‘event’?
- is dependency higher in the extreme tail of events, or is it linear?

A range of techniques can be used for building dependency between different risks into a capital model. In practice, the methods available will depend on the form of the assumptions for the risks in question.

A distinction can be made between methods that produce occurrences affecting more than one risk, and those that correlate the losses across the risks without generating underlying common catastrophes.

7.5.2 Use of event sets and event ID mapping

Event sets provide the simplest way of modelling dependency between risks, as long as those are modelled using the same event set. Basically, losses are accumulated for each event, and the dependency is automatically captured.

Where the risks being linked are from different event-based models, it may be possible to introduce dependency by mapping the event IDs between the different event sets using physical considerations. This mapping may only be partial as it may not be the case that an actual loss event will affect all risks at the same time. This process might require the splitting of individual events in the outputs, and their rates, into two separate events. For example, not all earthquakes lead to fire-following.

Companies need to take care to ensure the loss distributions of the risks on a standalone basis are not unintentionally affected by the mapping process. In some cases, a change may be necessary if the two models have incompatible assumptions. For example, if separate models are used for onshore and offshore wind exposures in the Gulf of Mexico, they may have different assumptions about the rate of Cat 5 events.

7.5.3 Using conditional probability and correlation matrices / copulas

Where the underlying models to be correlated are based on frequency and severity distributions, a different approach is needed. Catastrophe model outputs have an inherent, peril-determined dependency structure which needs to be simulated in some way.

A common approach is to use conditional probability to deal with the correlation in frequency. For each simulation, this method determines which risks (or peril regions) are affected based on the conditional probability distributions, and then samples the severities appropriately.

Where more than one risk is affected, a copula can be used to specify the dependency between the severity distributions for the affected risks. The greatest challenge is to create the correlation matrices and choose appropriate copulas, as these methods are often used when data is sparse.

7.5.4 Conditioning distributions

This is an approach that can be used to correlate losses from the correlating risk when these are considered separate events or otherwise cannot be accumulated for reinsurance purposes. For example, some climate oscillations are believed to affect the frequency of perils in widely-separated parts of the world such as US hurricane and Australian bushfire.

Here, a new variable could be introduced as a mixture parameter, based on a climate index, to condition the separate distributions. The mixture parameter could be discrete; for example representing “low” or “high” activity years or continuous.

For each simulation in the model, the sampled mixture parameter or state can be used to adjust the underlying frequency distributions - or event rates if using outputs - before they themselves are sampled.

Again, companies should take care not to introduce unintentional changes in the individual risk results. For example, a gamma mixture distribution applied to Poisson frequencies will result in the overall frequency being negative binomial.

7.5.5 Aggregate loss dependency

Where some dependency is considered important, but the previous methods cannot reasonably be applied due to lack of data or materiality considerations, an alternative is to 'retro-fit' dependency into the model after the separate risks have been simulated independently.

This is often easiest to implement using simulation shuffling, that is: to renumber the simulations to achieve the desired rank dependency structure between the annual aggregate losses of the two risks. For example, a modeller could sample correlated standard normal distributions using a preferred copula and parameters. Then, for each risk, he/she could rank these samples and replace the ranks with the corresponding loss for that rank from the individual independent simulations.

This cannot produce common cat events, but it does correlate the annual aggregate losses. It also preserves the distribution of the individual risks.

7.5.6 Implicit dependency

Where the quantification of a NMR is in the form of a probabilistic output adjustment, the dependency is implicit in the adjustment, whether that adjustment is supplied separately or already applied. In this case, no specific action on dependency needs to be taken in the capital model. For example, a flat 10% uplift to a loss output to allow for an un-modelled coverage is implicitly 100% correlated with the original modelled coverages.

7.5.7 Perfect correlation

In many cases, given the materiality of the risks in the overall model, detailed investigation of the dependency and subsequent simulation may not be worthwhile. If so, it is not uncommon to make the pragmatic - and cautious - assumption of perfect correlation between the risks. This is easy to implement using most of the techniques above.

7.6 Testing the model

7.6.1 Back testing

One important aspect when integrating non-modelled perils into the capital model is to sense-check the results from the overall capital model once the new perils have been included.

Several methods are commonly used at this stage and some examples are given below.

7.6.2 Sense check of expected losses

NMR captured in the capital model should go through the same testing process as all other parts of the capital model, depending on materiality.

The notional upper bound of the distribution used for NMR is given by the total exposed limit (or aggregate) for the portfolio/region/peril being considered. The return period of event at which this upper bound could be reached must be assessed using expert judgment, drawing on both underwriter experience and the purchasing behaviour patterns in a given market, coupled with the catastrophe specialist's knowledge of the peril(s) under consideration.

The area under the exceedance-probability curve for the risk is the annual expected loss for the portfolio/region in question, and it can be calibrated to correspond to the overall sum of cat premiums allocated to the region/peril, and recorded in the company pricing or contract management system.

The lower bound, (which may be zero), is the minimum return period at which losses are expected to arise, or below which natural catastrophe losses may be modelled as 'attritional'. This is an important parameter for the shape of the distribution.

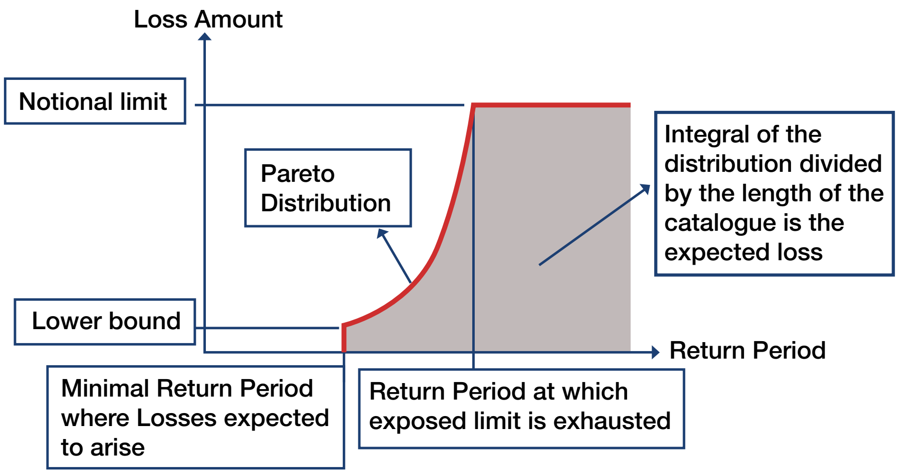

Figure 14: Fitting a distribution using total limits and cat premium

A first step to test the output from the capital model would, therefore, be to compare expected annual loss, maximum loss, and attritional use loss against expectations, which might be based on expert judgement, past experience, or other sources.

If the output does not match expectations, companies should investigate the reasons. Past experience might just have been very fortunate, the business mix may have changed, or the implementation or quantification was incorrect.

7.6.3 Sense check number of overall cat losses and severity of individual events

In addition to the overall expected loss, it may be useful to check how many catastrophe events are generated per simulation year, with and without NMR. The company should ask if the number of events, both the average and for extreme years, are sensible and within expected range.

If possible, the final severity of individual non-modelled events can also be sense checked. A simple, high-level check of the output can reveal issues in the financial structures applied or dependencies inside the model.

7.6.4 General items to check against

In addition to comparing the expectations, it might be useful to test:

- do results stack up? Does the sum of individually modelled loss components equal the total, gross and net?
- does the correlation work as intended?

These two checks use expert judgement to understand if the output from the capital model is as expected. If the results are significantly different from what would have been estimated, further investigation may be needed.

7.6.5 Back testing with past experience

Net loss ratios, produced by the model at divisional or entity level, are a good starting point to test the reliability of the results. When comparing losses that have occurred with modelled numbers, it is important to account for business growth, especially in emerging risk areas where no prior history exists. It is also important to understand that historical data will usually only go back a relatively short period of time and hence will not necessarily be a valid test for tail events.

Nevertheless, it would be sensible to compare, for example, the loss ratios over the last decade with where the individual years would appear on the current net loss ratio exceedance-probability curve to identify obvious outliers.

In addition, incorporating non-modelled perils can increase the average annual loss significantly, especially if double counting with attritional losses has occurred. These sorts of issues can be identified by comparing expected annual profit for the overall company, with the model output.

7.6.6 Scenario analysis

Companies should choose specific scenarios which include NMR to understand how these losses feed through the model, and to assess if all correlations as well as diversification and risk transfer options have been taken into account. For example, certain outwards reinsurance structures may have been missed for the new perils.

EXAMPLE 5 - OFFSHORE PLATFORM

A company is modelling on-shore losses to US hurricanes using a probabilistic vendor model. Losses to platforms in the Gulf of Mexico, however, are estimated using a loss distribution based on historical loss experience. The distribution is then correlated to the vendor model losses by risk ranking, with a high on-shore loss resulting in the company choosing a point higher on the off-shore distribution.

Scenario testing is employed to validate the output, using a deterministic event for the Gulf of Mexico. It is found that reinsurance purchased against US hurricane losses has been applied before taking into account the off-shore portion of loss, hence not accounting for recoveries of losses to off-shore platforms.

7.6.7 Sensitivity testing

Sensitivity testing for NMR could include increasing the estimated losses of specific perils by a certain percentage, or increasing the dependency between perils. This will highlight if the company could be at a significant risk should losses have been understated, or linkages misunderstood.

7.6.8 Stress testing

Stress testing applies extreme changes to both scenario-analysis and sensitivity-testing. For example, a 1,000-year return period event could be chosen, or a chain of several catastrophic events in combination with an adverse investment market move.

This is especially useful for NMR due to the high uncertainty around tail events, and could, for example, highlight additional risks.

7.6.9 Reporting and governance

The results of the tests listed above should be included into the Risk Management control cycle described in Chapter 2. As discussed previously, it is important that concerns are discussed with the relevant experts in the catastrophe modelling or underwriting teams.

It is essential for good enterprise risk management practice that the testing of the capital model and its results do not just stay inside the actuarial or capital modelling group, but that the business has a chance to react, for example by considering appropriate mitigation.