Vignettes from the 2012 GIRO paper
Catastrophe Model Blending: Practicalities

Alan Calder, Andrew Couper, Joseph Lo[1]

Abstract

The use and scrutiny of multiple catastrophe models in risk management is commonplace in firms with significant exposures to natural catastrophes. Where a single model stands out as the most appropriate for the firm, it is natural to adopt it for accumulation management. Where two or more models are each judged to contribute significantly to quantifying exposures, blending needs to be an option. This is not a new concept: combining forecasts has been common practice inside and outside actuarial work, especially helpful for dealing with uncertain situations and models. This paper draws from the fuller 2012 GIRO paper by the same authors. It first demonstrates the uncertain nature of catastrophe exposures and modelling using the European Windstorm peril. It then briefly describes a technical approach to blending that is consistent between pricing and accumulation. As governance of the process is critical to the success of blending, the paper then engages in a detailed governance discussion.

Keywords Natural catastrophes, catastrophe models, combining forecasts, model blending, model governance, model validation

Consideration of multiple expert opinions, models and methods is standard actuarial practice when facing uncertainty in situation and uncertainty in method or model. Due to the uncertain nature of catastrophe risk, for firms with significant natural catastrophe exposures, pricing and accumulation are typically performed with views from as many different catastrophe models as practically possible.

Combination of reasonable forecasts is commonplace in many disciplines such as economics, medical sciences, weather and climate forecasting - (Armstrong, 2001) gives a good survey of how combining forecasts help to reduce error in uncertain situations. In general insurance, a popular terminology for combining catastrophe models is model blending. While reinsurers price catastrophe risk with blending multiple model outputs (e.g. by averaging resulting loss costs), it seems less commonplace for reinsurers to incorporate blending for accumulation management.

The present authors wrote and presented a full paper – *Catastrophe Model Blending: practicalities* (Calder, Couper, & Lo, 2012) – on this topic at last year's UK GIRO conference. We now present three vignettes from the paper for debate and discussions in this year's ASTIN colloquium.

- Section 1 considers European Windstorm as an example of a natural peril whose uncertain modelling motivates model blending. This section corresponds to Section 1.6 of the GIRO paper.

- Section 2 suggests a blending methodology that allows a consistent approach to both pricing and accumulation. This section is a very brief summary of Section 3 of the GIRO paper.

- Section 3 discusses risk management governance issues surrounding the use of blending. Ability to blend and adjust models is often viewed with suspicion, due to its associated power to alter model outputs with biased subjective judgements. However, as described in generalist forecasting papers such as the seminal (Bunn & Wright, 1991), the use of judgement in the context of statistical modelling can be helpful when structures (or, at least, audit trails) for the use of judgement are in

[1] Respectively, Group Head of Catastrophe Risk Management, Chief Actuary and Chief Risk Officer of Aspen Bermuda, and Head of Actuarial Research & Development. Corresponding author: Joseph Lo, London, U.K.; jo.lo@aspen.co

place. As governance is critical to the success of blending, we replicate the whole of Section 4 of the GIRO paper in this section, with minor editorial changes.

Section 4 concludes the paper.

1 European Windstorms

Significant catastrophe losses for Europe in the UK and France illustrate two significant issues of using catastrophe models for both pricing and for portfolio accumulations. The most significant peril from an insured loss perspective is generally windstorm, which are also referred to as extratropical cyclones as they form outside of the tropics in the mid-latitudes.

Two such large losses from Extratropical Windstorms in Europe in the 1990s are (estimates are taken from Table 9 of (Swiss Re, 2012)):

- Windstorm Daria in 1990 (LUNCO code: 90A) in the UK with a market loss for the UK of approximately US$8bn, from the UK and other parts of Europe, in 2011 values

- Windstorm Lothar in 1999 (Lloyds code: 99AD) approx. US$7.8bn in 2011 values for all of Europe including losses from, southern Germany and Switzerland

Such events spurred the development of models in Europe for windstorm. Several specialist model vendors and other market participants such as brokers developed catastrophe models to simulate losses from such events. Due to the lack of standard met office data and accepted methodologies for simulating European windstorms, models for European windstorm have tended to produce a significant variation in outputs on broadly similar exposures. Despite significant revisions and enhancements, a significant degree of variation in results can be obtained between models even at an industry loss level, using consistent original exposure assumptions:

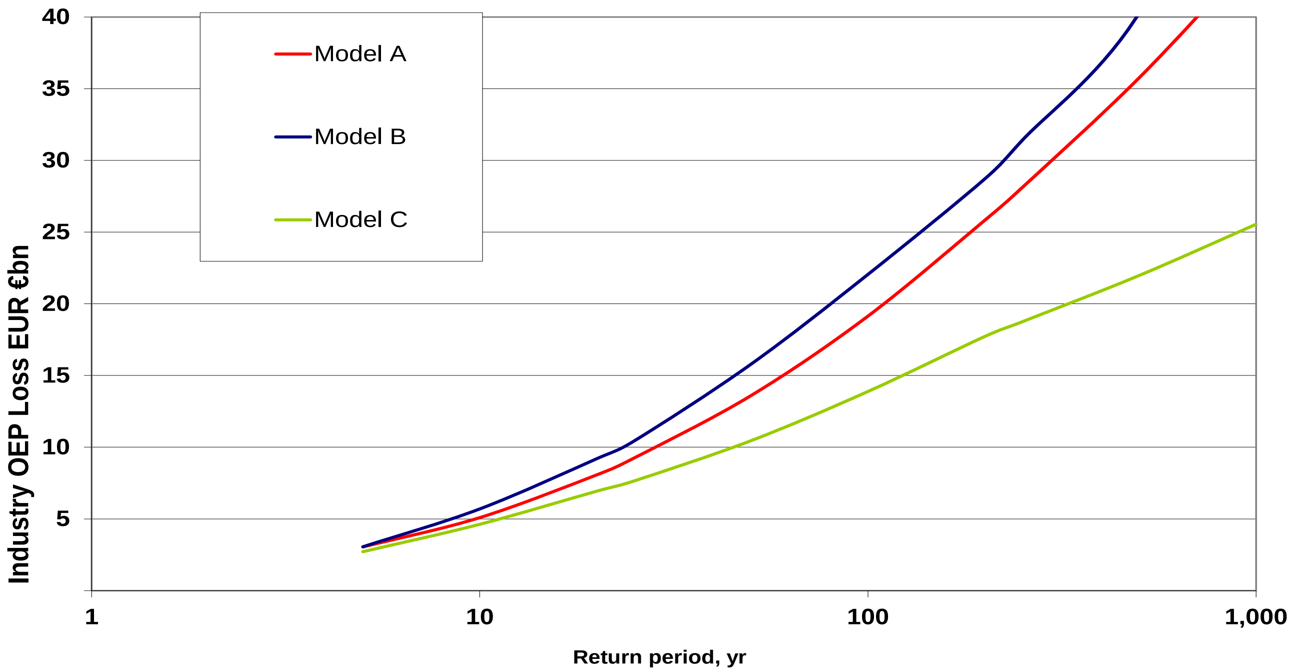

Source: Aspen model validation for European Windstorm model revisions, 2011

Secondly, further uncertainty in the use of such models can be demonstrated by benchmarking of various individual insurance portfolios in France, with recent loss events such as Windstorms Klaus in 2009 and Xynthia in 2010 where issues of trending actual losses to today's values are not as material as attempting to analyse more historical events. Although overall market losses are reasonably well simulated within 10% overall, individual portfolios in this example reveal differences of up to 40%-45%, in either direction:

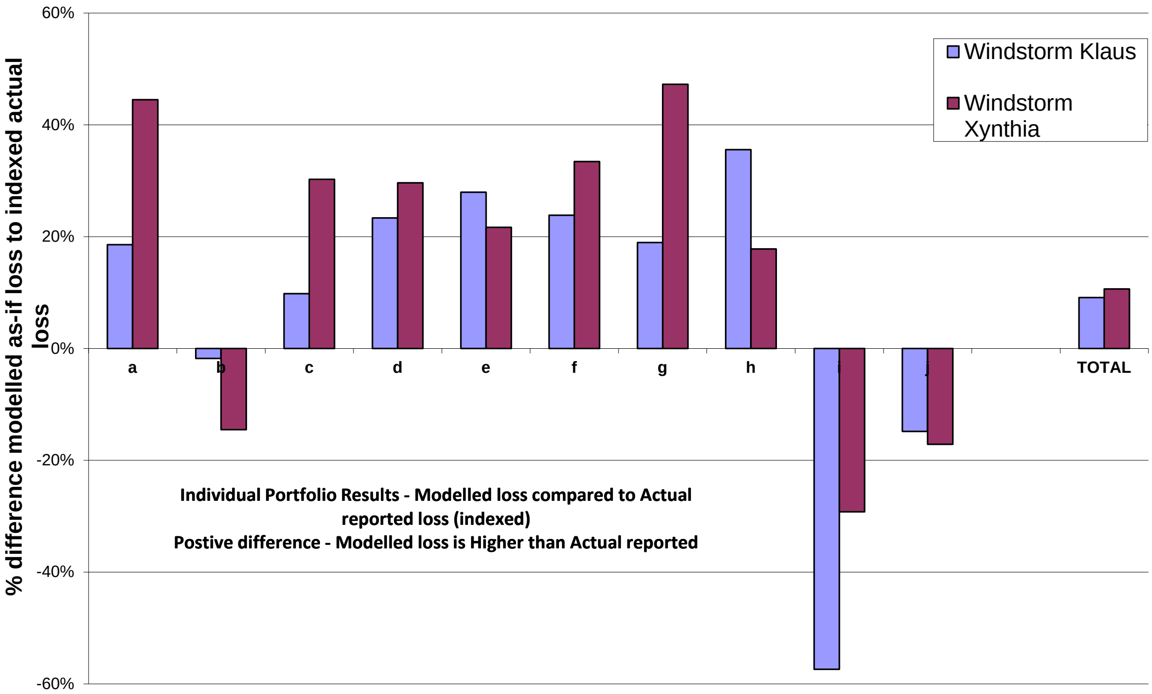

Source: Aspen model validation for European Windstorm model revisions, 2011

This principally reflects two main facets of catastrophe models. Models use "weighted average" assumptions for general vulnerability functions (which in Europe often tend to cover only broad risk types such as residential, commercial and agricultural), which may not always reflect the complete range of loss experience for specific types exposures written. Models typically use "weighted average" approaches when the user enters information without providing a high degree of granularity. This aspect of "averaging" can be reduced where the original data quality is improved through capturing more detailed attributes of the risks such as construction and year built, number of stories and roof types, etc. to reduce the levels of inference made by the model. However this requires the model developers to be able to discriminate by such characteristics: this may not always be the case, or can mainly be achieved from a significant degree of expert judgement rather than from actual loss experience. Even where such detail exists, a modelling company will derive a single mean damage relationship with a measure of uncertainty rather than produce a bespoke relationship for each and every exposure entered.

Particularly for smaller localised portfolios, such a portfolio may have significant biases in exposures on the ground relative to the geographical resolution of hazard within the model: essentially the resolution of the model may not be adequate enough to capture such effects. Finally claims management practices such as loss adjustment can also impact estimates significantly and again are "averaged" by current generation models. The above example demonstrates the variance of results possible within a single model – the use of multiple models may help to smooth individual biases, where the different models treat exposure and claims data differently.

Given the degree of uncertainty – at the overall exposure level and at the portfolio level – many brokers and (re)insurance companies have adopted a multiple model strategy in order to understand the significant variance in potential exposures. Indeed, it may be desirable to use a combination of actuarial models and catastrophe models, particularly for specialised exposures or to assess high frequency exposures. In our experience of using current generation catastrophe models, there can be a significant degree of overlap between what has been traditionally considered as "attritional" losses (modelled using an actuarial experience based model) and what was regarded as "large losses" (modelled using catastrophe models). In reality, small or medium sized storms are included within catastrophe models but results may not always align with trended experience resulting in further uncertainties for decision makers.

2 A Blending Mechanism

We now briefly describe a blending mechanism that consistently combines models at pricing and at accumulation. It is a two-step process: producing the *standard agreed blend* and then performing *bottom-up adjustments*.

The first step is to produce the *Standard Agreed Blend*. This can be considered as a top-down approach to blending, where the weights are pre-selected, and all risks in the accumulation are treated in the same manner and subject to these weights.

A “frequency blending” approach is employed in our example, based on simulations of years of catastrophe events. Such an approach has advantages – both mathematical and practical – over “severity blending”. A comparison of the two approaches is beyond the scope of this paper: the reader is referred to the original GIRO paper (Section 2) for a fuller exposition.

Following (Cook, 2011), pre-defined proportions of the simulated years will follow the component models, according to the pre-selected weights. All risks will follow these same pre-simulated years, so that accumulations could be done in a consistent basis.

The standard agreed blend is a useful framework that produces a blended event set – or indeed it is a catastrophe model in its own right. Such an event set forms the back bone for accumulating exposures. It allows consistencies (up to a point) with pricing, where the pricing blends the AALs with the same weights, and where the AALs are purely derived from the component models.

It is conceivable that underwriting and further R&D in the company would point to certain portfolio of risks to be treated differently from the top-down assumptions. A practical blending mechanism should allow for such deviations. We propose *bottom-up adjustments* as a set of tools for this.

These tools aim to adjust the severity of losses for each portfolio of risks. They do not alter the event set. This point is critical to the mechanism, since a key requirement for it is to be able to accumulate exposures. There are two at least two ways of doing so: scaling the severities uniformly and scaling the severities differently for different parts of the curve. Although the second way can be thought of as a generalisation of the first way, the first way allows underwriters to incorporate externally calculated AALs to be used for scaling.

The adjustments are performed on pre-treaty losses. More information is typically available from the original business’s perspective, rather than from the contract’s perspective. The market’s observed stability of stable reinsurance programmes notwithstanding, in theory, the contract and programme structures can change. It is more flexible to apply structures to the original business than to try and adjust the losses to the structures themselves.

Clearly, the adjustments apply more to reinsurance business than insurance business. However, it is possible that some of the techniques can also be applied to insurance at the portfolio level.

3 Governance

In order to make informed risk management decisions that rely on the use catastrophe models a sufficiently robust governance framework needs to be established. This should consider all aspects of the use of models including devising policies for selection of models, processes for the validation of models and oversight and review of implementation decisions. Sufficient evidence and documentation of the process is considered essential particularly given the requirements of the Solvency II regulatory regime regarding internal capital models which, for entities with significant catastrophe risk, rely heavily on the outputs of catastrophe models.

Various forms of governance can be considered. One more extreme model would be a strictly centralised approach mandating only the use of one single catastrophe model for all exposure classes underwritten (in both pricing and accumulation) with all operational aspects undertaken by a centralised risk management team. This has the advantage of absolute consistency of approach and simplicity of operation. The disadvantage of this approach is the risk of systemic errors in relying on a single perspective of a highly uncertain risk. This approach would neglect the potential benefits of using the best available models for differing peril zones, for example, or differing risk types.

At the other extreme would be a heavily decentralised approach with a central committee delegating the choices of model selection and operation to each specific underwriting unit freedom of choice to select and validate models. The advantage of this approach is that the most appropriate model for each class of business could be used but at the expense of consistency of approach. This would be a significant limitation for classes exposed to the same correlated geographical perils. Determining the appropriate level of correlations between business units becomes a more complex blending challenge particularly if no centralised risk management process exists. Also model updates and revisions become more complex to validate if a large and diverse number of models are deployed.

Various other forms of governance between the two extreme cases set out above are conceivable with: for example, a centralised approach to the selection and validation of models; but with some operational autonomy provided to decentralised users to determine the appropriate settings and assumptions for specific regions or business lines.

The remainder of this section is devoted to how model blending might impact governance of catastrophe model use in the company. We shall also highlight governance activities that we found particularly helpful. In particular, with the advent of Solvency II, various advice and guidance (for example, the ABI recently published a guide (Garnons-Williams & Zink, 2011) in this area) are in the public domain: our aim does not include a repeat of their content.

Just as our example blending mechanism is divided into two parts: the standard agreed blend and the bottom-up adjustments, we shall divide our discussion of governance into the governance of these two parts. However, the section should be read as a whole to gain a fuller feel of the discussion: the process presented has a circularity element to it, in which the information gained from the bottom-up adjusted blend would help inform (next iterations of) the standard agreed blend.

We start with the governance of the standard agreed blend.

3.1 Governance of the standard agreed blend

3.1.1 Vendor model validation

The process of setting the top-down view of risk should ideally be driven by a comprehensive in depth model evaluation process. This is often also referred to as model verification or validation process. Catastrophe model validation is both an art and a science, especially for modelling long return periods. The straightforward, sledge-hammer, way of validating models through comparing with experience data direct is impractical. Not only do we not have hundreds of years of loss data, in particular, to validate 1 in 200 figures and beyond, even if such data existed, the fast-paced socioeconomic flux of the world renders such data irrelevant very quickly. However, there are activities – which we shall share here – that we have found particularly useful to engage in to help the validation process.

This process begins with a full review of the model(s) being evaluated, including the key model assumptions and limitations. Often extensive documentation is provided by the model developers or vendors which can be the basis for an evaluation. All key parameters and data sources for each key module – hazard, vulnerability and loss/financial calculations – would be reviewed and understood.

However, as well as reviewing documentation, a more robust level of validation can be achieved by undertaking a rigorous and systematic series of model tests and review. We shall split this discussion into two parts:

- Validation of model outputs
- Validation of model assumptions

Expert practitioners carrying out this work can include atmospheric scientists and earthquake scientists within one's in-house Catastrophe R&D team. In addition senior catastrophe modelling specialists perform many of the sensitivity and parameter tests and extractions to complement the fundamental R&D.

3.1.1.1 Validation of model outputs

The validation of model outputs can be carried out with a series of reference exposures as a base. The use of industry portfolios is a common choice to provide a “weighted” outcome – the one that, we would imagine, model developers spend the most effort in. The use of significant portfolios, specific to the companies, is useful to gain helpful insights to judge the model’s suitability for use within the organisation. Many levels of output could be usefully assessed including:

- overall industry losses
- portfolio specific losses
- historical as-if losses
- losses by zone / sub zone
- losses by major risk types
- construction and occupancy classes

For high return periods, these outputs can then be compared and contrasted against the outputs of other models, in the context of the individual models’ assumptions and limitations. The rationale here is that independent efforts to model the same perils would help to highlight strengths and deficiencies of each other.

By “other models” here, we mean

- Current versions of available catastrophe models
- Previous versions of available catastrophe models (especially when we are validating a new version)
- Current version of the in-house standard agreed blend
- Current version of the bottom-up adjusted accumulations

A key feature here is that: just as blending would not remove all uncertainty for us due to, for example, the use of common historic hazard datasets, model assumption and output comparisons would not highlight all the limitations. The consideration of secondary events is helpful to supplement such comparisons (see below).

In addition, for lower return periods, if available to the company, the outputs can be compared against one’s own historic loss data and experience. Indexed historic market losses can also be helpful. We expect this to be familiar to actuaries – especially those working in the pricing area, where experience and exposure rating are often used side by side. An extensive example can be seen in a study done on the US Tornado Hail peril, (Smosna, 2012).

We also refer the reader to Section 1 where we have included two charts from a recent internal model validation exercise: one comparing the model outputs at an industry level, and the other comparing historical losses against model outputs, at a portfolio level.

It goes without saying that the opinion of experienced underwriters is helpful on these outputs. With the example blending mechanism, their opinion now flows through quantitatively in the bottom-up adjusted accumulations (see Section 3.2.2 below), as well as via informal conversations and formal discussions in governance committees. Qualitative discussions can also assist this process: for example, Q&A with the model developers or other market participants such as broker evaluation teams.

We have found it useful to cover the modelling of secondary events following the modelled events in this assessment, particularly to determine the extent of the limitations of using a given model. Examples include:

- tsunamis following earthquakes

- fire following earthquakes
- storm surges from windstorms
- construction material and labour cost inflations due to sudden post event surge of their demand
- the pertaining insurance environment – for example, legal interpretations of Terms & Conditions

3.1.1.2 Validation of model assumptions

Where made available, individual key parameters can be assessed. We give two examples here: one on the hazard side, and the other on the vulnerability side.

An example for hazard parameters for earthquake would be the range of event magnitudes, considered in conjunction with their associated frequencies. Ideally hazard parameters should be benchmarked against independent scientific data. For example, rates of modelled event recurrence can be compared to recognised earthquake catalogues with published levels of completeness. Should a model noticeably diverge against accepted scientific data, further analysis maybe required to establish the cause. An example may be that a model developer chooses to add additional caution for extreme events. Where assumptions agree this also helps to demonstrate a level of confidence and understanding in a given model rather than simply relying on the model to have reasonable assumptions.

Similarly, independent sources of loss or vulnerability data could be used to review and evaluate models to give evidence on their fitness for purposes. For example, historical scenario results can be compared to observed losses to determine if any biases or limitations exist. This is often best done against a representative sample of portfolios. Alternatively if no direct claims experience is available in sufficient detail, engineering based data or damage surveys can be used to assess vulnerability curves or at least relative views of vulnerabilities, such as residential compared to commercial risk.

3.1.2 Governance committees and model adoption

Once a model evaluation is finalised, for further discussion and review, written findings could then be presented to any governance committees or processes, established with a specific remit to assess catastrophe risk. The written report should ideally contain, as well as the key findings, recommendations for the adoption of the given model(s), possibly with adjustments. In situations where multiple models are considered, the strengths and weaknesses of each model in the key sub modules (hazard, vulnerability, loss calculations) are given. Should no clear "winner" emerge from the technical assessment and the models concerned have sufficient credibility; a case for blending models is considered an appropriate response, particularly for regions with significant exposures. Consideration of the weight of evidence will be undertaken to devise the appropriate weights for "blending" two or more models.

Thus any blend of models can comprise:

- 100% of unadjusted view of a single model if it is considered robust enough, or
- an adjusted view of a single model, or
- a blended view of models, also possibly with underlying adjustments

We recommend the make-up of governance committees determining any such blend should comprise relevant stakeholders with sufficient experience of catastrophe perils, with a careful balance of underwriting and risk management functions to ensure appropriate oversight. Decisions could be taken on a "materiality" basis as a guiding principle based on the undertaking's exposures. For example a more senior committee can be formed responsible for governance of the model selection and blending weights for major peak zone perils, with scope to form a more junior committee responsible for decisions regarding less material perils for the (re)insurance undertaking.

Finally, in agreement with Chapter 5 of (Garnons-Williams & Zink, 2011), we find documentation of model selection decisions and the rationale behind them to be helpful: not only does it provide audit trails, but such documentation encourages model selection in a disciplined manner.

3.2 Governance of the bottom-up adjusted blend

The bottom-up adjusted approach allows underwriters the freedom to adjust the default blends on an account by account basis for any of the modelled catastrophe perils and zones. As mentioned in Section 2, two main types of adjustment are enabled, both of which adjust the pre-treaty simulated losses. The treaty terms are then applied to the adjusted losses. The first method allows underwriters to enter a revised expected loss to a layer by either reweighting the constituent modelled results, for example, shifting from a 50:50 split of given modelled outputs, to favour a specific model, such as to use a 60:40 split. Alternatively, the expected loss adjustment by layer maybe sourced from an actuarial based model providing a burning cost output. The second adjustment method is to adjust the exceedance curve itself by adjusting at different loss percentile points. For example losses can be adjusted at the 1 in 10, 25, 50 years, including non-linear type adjustments.

Ideally the underwriter would be required to document the basis of the adjustment and could also choose one of a few predefined types of adjustment to provide clarity on the basis of adjustment.

Overall two distinct levels of governance deployed around the use of the “bottom-up” view:

- “Day to Day” via various levels of peer review for individual uses at account level
- Strategic monitoring and review of the overall use, impact and extent on the entire portfolio by comparison of the aggregated result to the standard blend

3.2.1 Account level governance

The governance of the adjustment process at an account level can be achieved well by three levels of review and sign off, calling on the expertise of senior underwriters, catastrophe risk modellers and actuaries. The first two are part of a broader peer review framework designed to consider all elements of the underwriting process. Having staff with the requisite skill sets and experience to properly assess a complex problem and set of outputs is critical to providing effective peer review.

- The first process is the underwriting peer review process. Here, a more senior underwriter is engaged to review and sign off all aspects of the underwriting including any pricing assumptions. This process is undertaken in conjunction with company’s approved pricing policies and processes, and the catastrophe risk modelling process, which is itself subject to peer review for material accounts “pre-binding” of business.
- Secondly, all bound accounts would be subject to a Catastrophe Risk Modelling team review undertaken by a modeller independent to the original modeller participating on the account. The governance of such reviews can be directed to ensure accounts with the most significant exposures are prioritised and reviewed by the most senior analysts. The review should challenge the basis of any adjustments considered to lack robustness.
- The third level of review is targeted purely at specific accounts: these could be carried out by a pricing actuary. The need for the review is triggered by the materiality of the adjustment (either positively or negative at certain threshold, which can be adjusted centrally), or if the exposure commitment exceeds a certain materiality threshold (again defined centrally and applying to all underwriting units). Furthermore, to cover the remaining contracts, a random sample approach could be introduced forcing a certain percentage of accounts to have an actuarial review.

Looking at publicly available literature, discussions on the subject of peer reviews are surprisingly limited, even though there is widespread agreement that peer reviews should be part of actuarial work in general. A comprehensive discussion is (Kucera & Sutter, 2007), and we also refer the reader to (Gibson, 2008) for an account in the context of reserving. The few detailed discussions on peer reviews usually centre on reserving, and we have yet to find one engaging in the subject of peer reviews on the use of catastrophe models. We now outline a few thoughts for future research.

The aims of peer reviews are likely to be similar between reserving and catastrophe modelling (e.g. ensuring processes are appropriate, spotting material errors, improving methodologies, helping with appropriate communication of results). However, peer reviews in the context of catastrophe modelling can require different approaches.

- The day-to-day multi-disciplinary nature of catastrophe modelling requires different communication techniques than an actuarial peer review of a reserving report produced by another actuary.
- Since we can be considering tail events, the data is likely to be much sparser or non-existent at these levels for validation and calibration.
- Probabilities ("1 in 100") are the more usual statistics than the mean ("Best Estimates"). Together with the previous point, similarities with peer reviewing capital model outputs are apparent.

In light of these observations, future research on peer reviews on the use of catastrophe models would benefit from considering latest practitioner insights from the compliance of Solvency II's validation test (see, for example, (Chhabra, Validating internal models: a practitioner's perspective, 2011)). Behavioural aspects of communication and eliciting expert judgement on remote return period events will also be important: see, for example (Chhabra & Parodi, Dealing with sparse data, 2010) and (Arbenz & Canestrano, 2012). The work of Daniel Kahneman and Amos Tversky (the former a 2002 Nobel prize winner in Economic Sciences) are likely to be helpful: Kahneman's recent book (Kahneman, 2011) has attracted good reviews.

3.2.2 Portfolio level governance

For major exposures, regular quarterly reporting of overall accumulations for the entire portfolio incorporates comparative analytics of the top-down standard agreed blend outputs and the bottom-up view. Significant variances between the two bases of reporting would be detected. Should this cross a materiality threshold, further analysis and report can be undertaken to investigate if any systematic pattern is applied across all accounts within a given accumulation zone, or if the adjustments are restricted to individual cases. Depending on the nature of the adjustments applied, and the overall impact, scope then exists to review in further depth within any committees mandated with monitoring and controlling overall catastrophe accumulations risk assessment. These committees can be identical to those alluded to in Section 3.1.2.

We expect any relevant governance committees comprising various stakeholders directly involved with catastrophe risk assessment but also those independent of the underwriting process in Risk Management functions can provide oversight. These committees would determine if the basis of the adjustments provides sufficient indication that further research projects should be commissioned. Such projects would aim to independently support adjustments with an eye to then "re-set" the top-down standard agreed view to correspond. However should the investigation instigated determine the basis for the adjustment to be insufficiently robust, the existing top-down view would then remain as the ultimate reference point for determining levels of catastrophe exposures.

4 Conclusions

The use of multiple catastrophe models is helpful for understanding catastrophe exposures against a company, as each model has its strengths and weaknesses. While model blending is commonplace in pricing catastrophe exposures, it is less so in accumulation management. Where one model stands out as superior, its selection for accumulation management is natural. However, when two or more candidates are judged to contribute significantly to quantification of accumulations, blending models must be an option.

There are two critical issues surrounding the blending of models: technical considerations and governance processes. We have described briefly a blending mechanism that allows accumulations and pricing to be consistently blended. The topic of governance was discussed in depth, from validation of catastrophe models and of judgemental adjustments, to governance structures. We linked the two issues by noting how, with the proposed blending mechanism, adjustments imposed by underwriters at pricing can be summarised back up for governance purposes.

5 Works Cited

Arbenz, P., & Canestrano, D. (2012, May). Estimating copulas for insurance from scarce observations, expert opinion and prior information: a Bayesian approach. *ASTIN Bulletin, 42/1*, 271-290.

Armstrong, J. S. (2001). Combining Forecasts. In J. S. Armstrong, *Principles of Forecasting: a handbook for researchers and practitioners* (pp. 417-440). Springer.

Bunn, D., & Wright, G. (1991). Interaction of judgemental and statistical forecasting methods: issues & analysis. *Management Sci.*, 501-518.

Calder, A., Couper, A., & Lo, J. (2012). Catastrophe Model Blending: practicalities. *GIRO Conference and Exhibition* (pp. 1-62). Brussels: The Actuarial Profession.

Chhabra, A. (2011). Validating internal models: a practitioner's perspective. Liverpool: UK Actuarial Profession.

Chhabra, A., & Parodi, P. (2010). Dealing with sparse data. *GIRO.* Liverpool: UK Actuarial Profession.

Cook, I. (2011). Using Multiple Catastrophe Models. *The latest issues surrounding catastrophe modelling* (p. 38). London: The Actuarial Profession.

Garnons-Williams, T., & Zink, U. (. (2011). *Industry Good Practice for Catastrophe Modelling: a guide to managing catastrophe models as part of an internal model under Solvency II.* ABI.

Gibson, L. (2008). Defence against the dark arts: a practical handbook for reserving actuaries. *GIRO.* UK Actuarial Profession.

Kahneman, D. (2011). *Thinking, fast and slow.* MacMillan.

Kucera, J., & Sutter, R. (2007). Peer Review. *CAS Spring Meeting.* Orlando: CAS.

Smosna, H. (2012). Tornado / Hail: To model or not to model. *CARe.* Boston: CAS.

Swiss Re. (2012). *Natural catastrophes and man-made disasters in 2011: historic losses surface from record earthquakes and floods.* Swiss Re.