



Internal Model Industry Forum:
Diversification benefit:
understanding its drivers and
building trust in the numbers

Executive Summary

Internal
Model
Industry
Forum

Supported by:



Contents

Foreword	1
Introduction	2
The new challenge	3
Concentrate on getting the dependency framework working	4
Dependency structure	5
Causal drivers – what are they and how does the process work?	6
Statistical assessment of dependencies – what are they and how are they used?	8
Validation of dependencies	9
Summary	10
Appendix A – Case study on determining a dependency structure	11
Appendix B – Detailed description of using causal analysis to assess and validate dependencies	13
Appendix C – Linear Correlation and the SCR	16
Appendix D – Calibrating dependencies	19
Appendix E – Validation tests for dependencies	23
Our project team	28
Our steering committee	28
The Internal Model Industry Forum	29
Who are the IRM?	29

©2015 The Institute of Risk Management.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise without the express permission of the copyright owner. Permission will generally be granted for use of the material from this document on condition that the source is clearly credited as being the Institute of Risk Management.

IRM does not necessarily endorse the views expressed or products described by individual authors within this document.

Foreword



The Solvency II legislation recognises that the overall risk exposure of insurers can be reduced by the diversity of their business. This is because the adverse outcome from one set of risks can be offset by a more positive outcome from a different, uncorrelated, set. The amount of money at stake is significant: this ‘diversification benefit’ could potentially reduce the capital that an insurer is required to hold by up to 50 % , which could amount to billions of pounds for a major multiline insurer.

However, understanding and calculating the drivers of diversity within the internal model, working out the appropriate level of diversification benefits and making sure that this results in a credible number for the firm and the regulator are not easy tasks.

This booklet looks at how insurers’ internal risk models should take account of the potential for diversification credit and offers some suggestions as to how modelling assumptions can be validated and communicated in order to produce a result that the organisation and the regulators can understand and trust. It is one of a series being produced by the Internal Model Industry Forum (IMIF) offering guidance on the validation, communication and use of insurers’ internal risk models in order to create real value for the business.

I would like to thank the members of the IMIF workstream who produced this work, and particularly workstream leader Justin Skinner from QBE and supporting consultants Milliman, for their work researching and developing the approach in this booklet. The members of our IMIF Steering Committee also provided valuable overall project guidance and quality control. We are also grateful to representatives from the Bank of England (PRA) who have enabled us to maintain a continuous and positive dialogue between industry and the regulator on the work of the IMIF.

I would also like to thank our sponsors Milliman, PwC, EY, Deloitte, and KPMG. Also, thanks are due to the Institute and Faculty of Actuaries and to ORIC International and other IRM practitioners for their input to this project. As a not-for-profit organisation, IRM relies on enlightened industry support to help us publish documents like this. It is this kind of support that helps help us maximise our investment in the development and delivery of world class risk management education and professional development.

**Jose Morago,
IRM Chairman and Founder of the
Internal Model Industry Forum (IMIF)**

Introduction

At its most simple, diversification benefit arises when two processes are not completely dependent on each other, and a bad (good) outcome for one process does not necessarily mean a bad (good) outcome for the other. Dependency and diversification are opposite sides of the same coin; when the strength of a dependency is increased, the level of diversification benefit is reduced.

Diversification benefit within internal models arises from a number of places, whether it is between risk types (e.g. between insurance risk and market risk), within a risk type (e.g. between casualty and motor underwriting risk) or through a common driver (e.g. inflation and interest rates). At each stage when there is a diversification benefit, the board and senior management need to be given confidence that the numbers are appropriate. But how can this be achieved? It's difficult to argue that risks operate completely independently in the extreme scenarios adopted within the internal model, but diversification between risks is the fundamental premise of insurance. For many insurers, aggregate diversification effects will be the most significant determinant of required capital, with a 40% to 60% reduction in the capital required between risk types not being atypical for a large well diversified insurer or reinsurer.

But modelling dependencies and setting appropriate assumptions is far from easy. Data is typically available for market risks, but the choice of data source and time period can have a significant impact on the answer. Insurance risks, on the other hand, often have insufficient data to produce any credible statistical results. Overall, a significant reliance must be placed upon expert judgement.

This booklet aims to support the analysis of diversification benefit, from the start of the diversification journey, and working its way through to the validation of the outcomes. It is by no means a definitive guide and cannot guarantee your diversification benefit calculation is appropriate (particularly since no model is right, but some are useful), but will offer a few practical pointers from the combined practical experience of the working group and wider Internal Model Industry Forum (IMIF) members.

In the next pages (and supporting appendices) we:

- explore how to get a **dependency framework that is working** reflecting the risk profile of the firm
- answer the question “what is meant by correlation” by looking at **causal dependencies** that help identify and understand a common driver of a relationship between two items of interest
- reflect on the important features of the relationship between two or several variables to capture by introducing **statistical dependencies**, a modelling technique that helps produce the desired joint behaviour
- explain the different tools available to **validate dependencies**

The new challenge

Setting parameters requires data to obtain credible estimates.

You need the right quantity (usually the more the better) and quality (accurate and appropriate) of data. Think about assessing the combined annual variability of a combined motor and casualty book of business and how much data is required to support it.

- Expected loss experience – 7 to 10 years of history for both accounts should give you a good understanding provided the book is relatively stable.
- Variabilities (e.g. standard deviations) – 10 to 20 years of history will give you a good indication.
- “Average” correlation – 30 to 60 years of combined data to provide a credible estimate.
- “Extreme” correlation or tail dependence, you need even more data. And how many years of data do you really need to credibly calibrate a 1 in 200 year event?

But how many companies have 30 to 60 years of good quality historical data? And of those companies, how many do you think believe their accounts have been stable enough to be fully representative of the future? And beyond this, how many people believe that drivers of dependencies (such as inflation) will look the same in the future as they have for the past 30 years? Or 200 years? This potentially calls for cautious expert judgments.

We haven't carried out any surveys, but our guess is that the answer is none. So, how do you solve an insolvable problem? The answer is to reframe the question so that the process of setting dependencies aids validation, risk management and insight into the business.

Although the following covers dependency modelling and validation in detail, all users and validators of dependencies in the internal model should make sure they focus on:

- Model capture: whether the model captures dependencies regarded as key by the senior management team and those that have the most impact on the outputs of the model for its intended uses;
- Model extent: the extent to which modelled dependency risks reflect severe but plausible assumptions that are of the nature the board wishes to affirm in its risk assessment and risk appetite considerations.

Validation tests like sensitivity testing and cross examination with relevant scenarios are great tools for drilling into the internal model to find out what really moves the dial in respect of capital.

“Focus on what is important in the internal model, and don't drown in the detail.”

Concentrate on getting the dependency framework working

When there is not enough data available in an internal model, expert judgement is often used to supplement it. However the input from experts has to be carried out with fairness, balance, and most importantly of all, transparency.

The following shows the process to work through with dependencies in internal models, at each step of the process the validation should be conducted and consolidated at the end of the process.



- Where would you expect relationships to exist?
- Make sure you can incorporate them.



- What relationships do you believe exist within your business?
- If you can identify the driver of these relationships then you can incorporate that relationship directly into the model.



- Apply a statistical relationship between classes directly rather than using a causal driver.
- This can capture “background noise” between classes to incorporate relationships where you just can’t identify a specific driver.



- Test what you have done to see how well it works
- Back-testing and simpler aggregated models excel here
- Understand which dependencies have not been incorporated from the methodology and their materiality, these should be recorded as limitations.

Dependency structure

The choice of dependency structure of an internal model can have a significant impact on its functionality, run-time and results. Once chosen, it is difficult to change. It is therefore important to invest time in determining it before the model build begins.

The important considerations are:

- **Senior management.** To gain buy-in, it is crucial that the model captures the key dependencies that senior management identify. An example would be the relationship between catastrophe losses and reinsurer default.
- **Market standards.** You will have to justify your choice of dependency structure to the stakeholders. Missing out a widely used dependency could raise questions with stakeholders, resulting in significant additional work and, potentially, capital loadings.
- **Balancing granularity with complexity.** The model should be granular enough to represent the behaviour under consideration sufficiently for the intended use and gain acceptance within the business. However, this should be balanced with the realisation that almost all dependency choices will be a matter of ‘expert judgment’, with the associated overhead under Solvency II. Furthermore, too much granularity can lead to spurious accuracy, obscure the fundamental drivers and lead to model instability.
- **The intended uses of the model.** A dependency applied at a high level may be an adequate reflection of a process at the aggregate level, however if the model is used for more granular reporting, the dependency may not reflect the process. For example, the value of a reinsurance contract can be strongly affected by the assumed dependency strength between the lines of business that it covers.

Appendix A shows a simple case study on determining a dependency structure within a company.

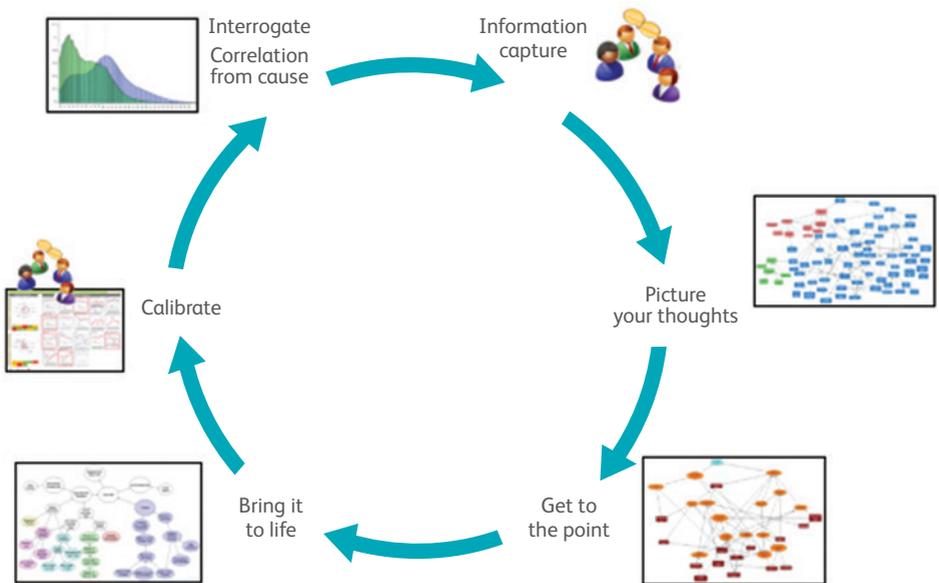
“Invest time upfront to choose and understand the dependency structure.”

Causal drivers – what are they and how does the process work?

Causal dependencies look to identify and understand a common driver of a relationship between two items of interest.

For example, an unexpected increase in inflation is likely to drive increased claims liabilities if they are inflation linked, and excess investment returns since interest rates are likely to increase. In this example, inflation is the cause of the dependency.

Following is a diagrammatic representation of how causal dependencies are incorporated into an internal model.



“The key is to ensure that expert judgements are brought together in a consistent way and articulated within a framework which can facilitate the communication of results, validation and delivery of additional business value from the exercise.”

The aim is to determine a structure describing how the outputs arise through the interactions of the inputs. It starts with brain-storming exercises to identify a detailed understanding of the causal relationships that could exist. The next step is to remove levels of detail (“complexity”) until the remaining structure tells us enough about the key relationships we have to consider without clouding our understanding with too much detail (a “minimally complex” summary). These relationships can then be modelled in a way that joins inputs to outputs transparently, and also that enables us to be very clear about where we have uncertainty. Scenario analysis and back-testing can be used to explore the dynamics of the model and the associated implied correlations under a wide range of conditions.

There are a number of ways of identifying and incorporating them into an internal model, and Appendix B works through a case-study that has been carried out for a life insurance company. It is worth noting that this is at the best practice end of causal analysis and simplifications can be made.

Statistical assessment of dependencies – what are they and how are they used?

Statistical dependencies can be modelled between two or more variables to produce a desired joint behaviour. This technique has its foundations in the same principle of understanding drivers of risk as causal modelling and using this as a basis for analysing the historical data and incorporating expert judgement.

It may not always be possible to build models with cause and effect relationships, so it is required in all other cases to apply statistical relationships between risks which have a likelihood of common occurrence. Having the right model structure at the appropriate level of granularity facilitates imposing dependencies in a meaningful way which can be effectively calibrated, challenged and validated.

The following list gives examples of variables to which copulas are often applied in General Insurance internal models:

- Reserve variability and premium risk variability within and between class
- Rate movements between classes
- Catastrophe claims between classes, both natural and man made
- Catastrophe risk and reinsurer defaults
- Default risk between reinsurers and brokers
- Asset returns, inflation and exchange rates
- Claims experience and market risk
- Operational risk and claims, or macro events

Appendix C shows the simple approaches adopted for the assessment of the capital requirement under the Solvency II standard formula (and other simple regulatory and rating agency models).

Appendix D discusses some of the ideas of how to calibrate statistical dependencies in an internal model.

“Statistical dependencies allow the easy inclusion of any relationship that you can think of, including elements of background noise to reflect the unknown unknowns.”

Validation of dependencies

Validation of dependencies is likely to be a multi-pronged approach with a wide variety of testing prepared and ready for review. Possible validation tests include:

- Sense check
- Independent assessment
- Stress and scenario testing
- Back-testing and aggregation testing
- Benchmarking – compare to the market
- Input / output checks
- Tail dependency tests
- Sensitivity testing

Appendix E includes more detail around each of the validation tests and how they can be used to assess the accuracy of dependencies in an internal model.

There is no single right answer so it's important to assess whether there has been sufficient challenge and review in a process which, by its nature, involves expert judgment. Having sufficient rationale for selections, evidence of testing and documentation to review is key, but, given this wealth of information, it is important for a validator to take a common sense approach knowing which areas are most material and need to be prioritised for checking.

It is worth focusing on the outcomes of the dependencies rather than the specific parameters themselves. For example, nobody really knows what a 30 % correlation between two classes of business really means. So it is better to look at the uncertainty of the aggregated net combined ratio rather than trying to validate 30 % specifically.

“Focus the validation efforts on the dependency assumptions that affect the outputs.”

Summary

Dependencies are the holy grail of internal models. Understanding the interactions between risks is central to the question of determining the company's capital requirements, as it's very often not one single risk occurrence that can hurt but (potentially) a correlated accumulation of many of them that causes the problems.

Regulators are keen that diversification benefits are fully justified, and potentially cautious given the parameter error in setting them, whereas companies will target what they consider to be the most realistic assumptions to ensure capital requirements are appropriate.

There is limited data available to set dependency structures and parameters, and this also means that there is a limited number of ways to validate them. It is critical that the dependency structure represents the company's best knowledge on the interaction between risks whether from underwriters, actuaries, senior management or the board. Having a clear and transparent validation structure means that users of the model will be better able to understand and challenge it and this tension between conflicting views on meaningful assumptions keeps the model real. It is for this reason that half of this booklet on validating dependencies has talked about how to set up dependency structures, how to look at causal models to identify drivers of dependency and model them, and finally how to think about using statistical dependency structures in internal models.

This booklet has brought together a number of internal model experts to try to shed some light on how they are used in internal models, and some ways that they can be validated. We hope that following it will enable internal modellers and Boards to know roughly which country the holy grail that is dependencies is in, if not the exact castle where it is hidden.

“Focus on the material dependencies, and ensure transparency at every step of the journey. This will give you, and your board, a fighting chance of understanding the drivers and building trust in the numbers.”

Appendix A Case study on determining a dependency structure

The purpose of this case study is to demonstrate a method for identifying the dependencies between risk types for the Internal Model.

The starting point for incorporating dependencies is getting the right structure (as outlined in the process diagram in the previous section). There are many ways of doing this, and workshops with relevant business experts are a good way of investigating this.

Meeting format

The workshop was set up as a two hour meeting of internal business experts, facilitated by the ERM team. Each member was asked in advance to think about how strong relationships could be between risk types.

Materials going into the meeting

- A simple grid for making notes. All green items needed consideration

	Earned Reserve risk	Unearned Reserve risk	Underwriting risk	Credit risk	Market risk	Operational risk	Group Risk
Earned reserve risk							
Unearned reserve risk							
Underwriting risk							
Credit risk							
Market risk							
Operational risk							
Group risk							

- An open mind from all participants
- A wider range of industry experience from the diverse range of business experts (covering Finance, Actuarial, Underwriting, Investments and Reinsurance) to provide different viewpoints.

Meeting discussion

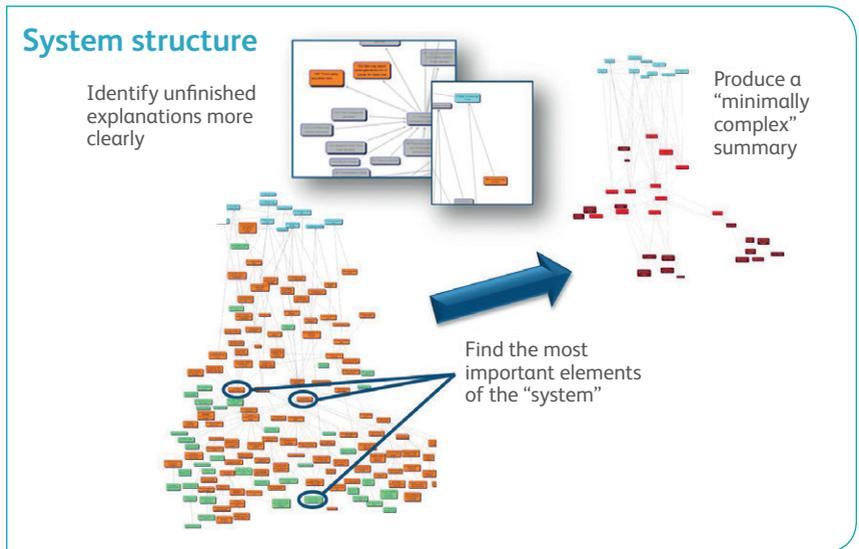
- Each cell was discussed individually, which captured a number of causal drivers of the dependencies. They were given a relative strength to prioritise future work on dependencies
- Benchmarks (from the standard formula and a friendly consultant) were reviewed to understand why the company might have stronger or weaker relationships than the benchmarks or why the benchmarks are not appropriate. The benchmarking material was only shared after the meeting discussion so that it didn't anchor people's thinking.

Any causal drivers identified during the individual cell discussion were cross-referenced at the end to see if they would impact other cells.

Meeting outputs

- A dependency structure based around the expert judgement of senior managers in the company
- A set of notes justifying what the dependency structure needs to incorporate.
- A first set of causal drivers to consider – see Appendix B for more on these.

Appendix B Detailed description of using causal analysis to assess and validate dependencies

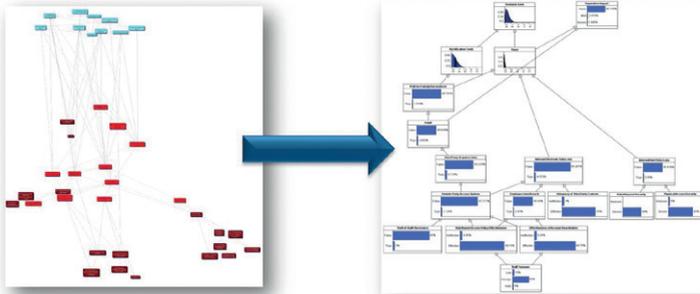


When experts provide narrative, they are explaining how a number of concepts relate to each other. We identify those individual concepts and link them in a map. As the various narrative threads are added, we build up the rich content that the expert is conveying. Colour coding helps visualise the data, highlighting risk drivers, controls and outcomes. For example, given the narrative “If policy-holder service standards deteriorate, this might directly lead to an increase in lapses. The decreased service standards may also lead to damage to company’s reputation and reduce its competitiveness in the market, both of which may also increase lapses.” Here, policy-holder service standards, damaged reputation and reduced competitiveness can all be linked up in the map as ultimately driving our output, lapse rates.

To reveal the minimally complex summary, we want to focus on the “most important” concepts. Whilst it is important to capture as full a picture as possible of the system dynamics, we can simplify things a little by analysing the connectedness of the concepts. It is important to validate the structure at this stage – both with those who contributed to the original discussions and more broadly with senior management. Does it look plausible and tie into your high level understanding?

“Mapping the system out makes it much easier to see where the common influences arise. The minimally complex view focuses on the key concepts.”

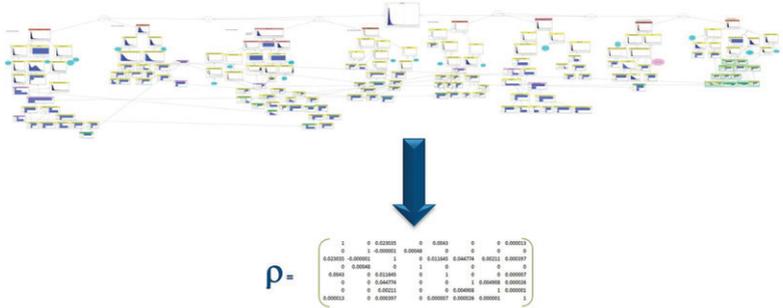
Bringing it to life



Now that the core features of the system are understood, it is possible to convert it into a causal model, with the minimal complex summary being the starting point. Expected outcomes and uncertainty around these outcomes can be developed based upon data and / or expert judgement, with the model ensuring that the various factors move together in the right way. Where two outcomes share an underlying driver in some way, their movements will display a level of positive or negative correlation. For example, where equity returns and lapse rates are both affected by interest rates, we would expect a change in interest rates to cause both lapse rates and equity returns to move at the same time.

“Causal modelling techniques can be used to simplify a complex business into relatively simple-looking models, which can be widely understood and easily explained.”

Correlation from cause



The correlation between risk drivers can be calculated directly within the model. They can be flexed through what-if analysis and back-testing to see how the correlation assumptions change under different conditions. This provides a highly intuitive and flexible approach to understanding and explaining the correlations.

Running the analysis in reverse can provide a powerful approach to validate correlation assumptions. Here we examine what the causal drivers would need to look like to generate a particular correlation assumption. This will help either lend support to the existing assumption or point towards a change in the risk driver assumption, or the causal relationships that have been incorporated. For example, we can fix the lapse rate output to be at a certain level and the model will allow us to see how the underlying drivers change. If doing this causes the drivers to display unusual movements or behaviours which are inconsistent with the chosen lapse output, then this may indicate that the relationships or calibrations of the model need to be reconsidered. If by setting lapse rates to be very high, for example, the model tells us that policy-holder satisfaction must also be high, then the calibration of the model may be incorrect.

“To the extent that two or more variables share underlying drivers, they are likely to exhibit some degree of correlation. This can be calculated directly from the model and explained.”

Appendix C

Linear Correlation and the SCR

To illustrate the concept of dependencies and the impact that they have on diversification benefits which can arise, a simple example is presented using linear correlations. A presentation on the strengths and limitations of this approach will also help act as an introduction to the rest of the paper which presents alternate techniques.

Background on linear correlations

There are a number of regulatory and rating agency capital models that use linear correlations to assess capital requirements. In a Solvency II context, the approach has been to stick with linear correlation as a simple tool but adjust the parameters to try to take account of its limitations. EIOPA made the following comment in their July 2014 paper on the underlying assumptions for the Standard Formula:

“Due to imperfections that are identified with this aggregation formula (e.g. cases of tail dependencies and skewed distributions) the correlation parameters are chosen in such a way as to achieve the best approximation of the 99.5% VaR for the overall (aggregated) capital requirement.”

Linear correlation is the approach adopted to aggregate risks for the purpose of the determining the Solvency Capital Requirement (SCR) under the Standard Formula in Solvency II. The basic formula is as follows:

$$SCR = \sqrt{\sum_{i,j} p_{ij} \cdot SCR_i \cdot SCR_j}$$

where SCR_i and SCR_j = the capital requirements for risks i or j , respectively and $p_{i,j}$ is the correlation parameter selected between risk i and risk j . i and j run from 1 to n , with $p_{i,j}$ being 1 wherever $i = j$.

Example: even risk capital split

Consider a simple example of two risks which generate required capital of £10m each. From the formula above, it is clear that the overall SCR will depend on the assumed level of dependence between the risks, as determined by the correlation parameter. The results for varying levels of dependence are illustrated in the table below:

$\rho_{x,y}$	SCR	
100 %	20	• $\rho_{x,y} = 1$: Risks are perfectly positively correlated, if one occurs so does the other. Thus we must hold the full £20m (£10m for risk i and £10m for risk j)
50 %	17	• $\rho_{x,y} = 0$: Risks are uncorrelated. There is therefore no need to hold the full £20m and we can hold a lower amount than this of £14m.
0 %	14	
-50 %	10	• $\rho_{x,y} = -1$: Risks are perfectly negatively correlated, one risk provides mitigation for the other.
-100 %	0	

The diversification benefits are clear; if the risks are deemed to have a correlation of 0 then capital required in this example is reduced by around 30 % (i.e. 1 - 14m / 20m). However, it should be noted that if one risk is dominant and has a higher capital requirement, then the diversification effect is reduced. For example, if the risks have required capital of £15m and £5m respectively, the required capital with no correlation falls by only 21 % rather than the 30 % when they are the same size.

Strengths and weaknesses of linear correlations

The use of similar methods in an internal model is acceptable, subject to the same caveat EIOPA made; the correlation selected should reflect the expected strength of the relationship in the tails. Such an approach also benefits from being transparent and easy to understand.

Advantages

- **Ease** – Results can be easily computed on standard software such as Microsoft Excel
- **Familiarity** – the logical step-on from the variance/covariance matrix approach of some companies' ICA approach
- **Transparency** – easy to see the impact of changes in the parameters, it's easy to communicate and easy for non-technical audiences to understand
- **Suitability** – Suitable where the marginal distributions of risk capital being aggregated are Normal and where the dependence relation between the risks being considered is linear i.e. the relationship does not depend on the area of the risk distribution being considered
- **Simplicity** – One set of subjective assumptions (correlations) rather than several subjective assumptions (degrees of freedom needed to implement tail dependencies using copulas)

Disadvantages

Unfortunately, for use in assessing risk based capital the approach can lead to over or under estimation of capital requirements for the following reasons:

- The shape of the marginal distributions can differ significantly from the Normal distribution – for example insurance risk distributions are typically quite skewed, such as catastrophe risk where there are some very extreme losses in the distribution.
- The dependency between the distributions is often not linear – for example there may be stronger dependence in extreme situations than for “average” outcomes.
- It does not capture causal relationships between variables so explanatory power may be limited

The second disadvantage can be managed through the selection of larger correlation factors which should reflect the expected strength of the relationship in the tails. However, this would mean that the overall probability distribution function generated by the internal model might not be appropriate for more day-to-day eventualities.

The other disadvantage of this approach is that the industry or market may be using more complex dependencies and therefore put a company using linear dependencies at a disadvantage.

“Using linear correlations can be a simple approach within an internal model to assess a point estimate, but they do so at a slight deficit to the use test.”

Appendix D

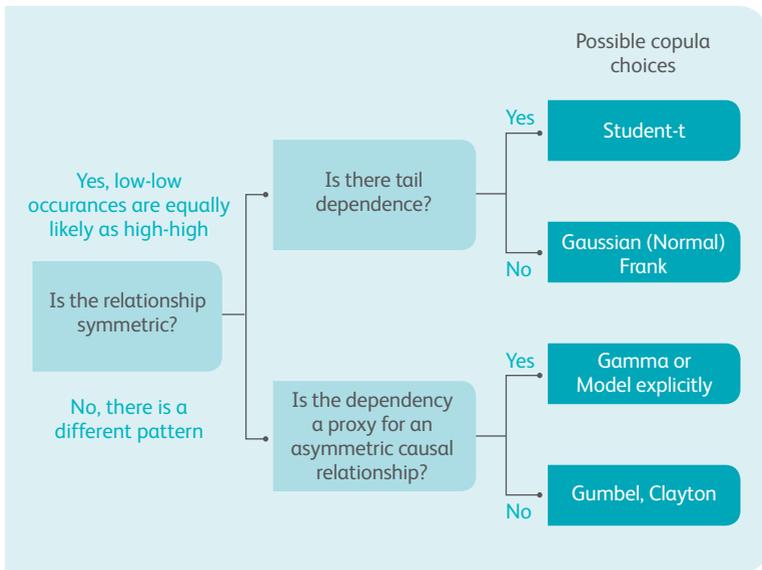
Calibrating dependencies

There are a number of considerations that you need to work through when calibrating statistical dependencies:

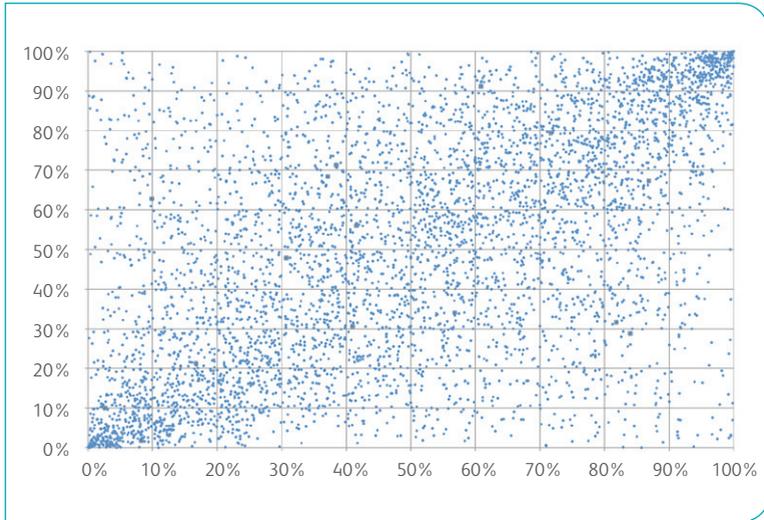
- **Data available:** Consider using company data, public data and market benchmarks. As an example, the difficulty in obtaining internal company data to measure extreme outcomes might be remediated by considering industry events or the latest scientific research.
- **Form of the distribution:** There is a wide variety of distributional relationships (also known as copulas) that can be chosen. The form adopted will very much depend on the expected features of the items being considered. For example, it is important to answer the question of whether the relationship becomes more significant in extreme outcomes, such as may be the case for reinsurer defaults occurring as a consequence of natural catastrophes.
- **Strength of the distribution:** Where historical data is insufficient to select parameters, a different approach needs to be adopted. This might require expressing the expected dependencies in terms of their risk drivers and using bandings to build up the relationship. For example, assessing the sensitivity of different lines of business which have common exposure to a particular financial risk variable, such as inflation, will help to indicate the lines of business that should have a stronger reserve risk relationship.
- **Expert judgements:** The selected form or strength of the dependency will always include an element of expert judgement, and this selection should be reinforced via workshops with underwriters, pricing and reserving actuaries, senior management and the board. Selections should be reinforced through the company's internal ERM practices such as risk identification exercises, risk monitoring and stress tests.

Form of the distribution

The most common form of applying a statistical distribution is to use a copula function. This booklet does not go into the science behind them because it is readily available elsewhere (e.g. <http://www.actuaries.org.uk/sites/all/files/documents/pdf/sm20100510.pdf>) and does not lend itself to too much simplification. However, following is a simple decision tree that companies have used to choose the copulas most commonly found in internal models. This simplifies a complicated subject into a small number of key decision points in a manner boards can readily understand.



Scatter plots can also be useful when considering which form of copula to use, and also as part of the validation of the choice. For example, the following scatter plot shows the use of a Student-t copula, which has a symmetric relationship with tail dependence (as can be seen by the clustering of dots in the bottom left and top right).



When choosing the form of the dependency, care should be taken to make sure there are no unforeseen consequences. For example, if you apply a two-way Gumbel dependency between A and B, and also between B and C, then the relationship between A and C tends to be quite weak. Is this appropriate? Should you be looking for a better relationship between A, B and C?

Strength of relationship

The calibration of statistical dependencies is reinforced by looking at key drivers of risk as this will help to analyse data, support expert judgements and increase the usefulness of results. Causal drivers are also a useful consideration here, but rather than being used to identify and model the root-cause factor, they can be used to gauge the relative strengths of relationships between risks.

Case study on getting a dependency structure

The following approach has been used to assess the strength of relationships between classes of business for reserves, premium rates, attritional losses and large claim frequencies.

- Identify a set of dependency drivers that you think might cause relationships between classes. For example, the following items might be identified as the main drivers of dependencies.

Frequency of claims		Severity of claims	
Driver	Description	Driver	Description
Economic Downturn	Increase of fraud, litigation against employers and claims are made that would usually be ignored	Tort changes	Increased severity of court awards due to the judicial environment
Equities	D&O more likely to be sued following drop	Inflation	Different lines will be susceptible to different forms of inflation, whether medical, wage or reconstruction cost inflation
Underwriting / pricing error	Poor risk selection leading to increased frequency	Underwriting / pricing error	Common poor risk selection leading to increased severity
		Claims handling	Common poor claims handling leading to increased severity

- Categorise (e.g. high/medium/low) the impact of each of these factors on each class of business for each type of dependency being calibrated. For example, both employer's liability and public liability will have a high susceptibility to tort changes.
- Assign each category a strength of relationship. This is driven by expert judgement, and can be sensitivity tested after the initial selection. For example, a high tort impact might drive a 20% correlation.
- Build up an aggregate correlation between classes based on these relationships. So for the EL / PL example, there would be 20% from them both having high susceptibility to tort changes. This then gets increased if they also had a high susceptibility to inflation (which they do).

This type of dependency decomposition gives a simple framework for assessing the strengths of relationships, although it does rely on expert judgement which needs to be validated.

“Consistency and clarity of approach is key.”

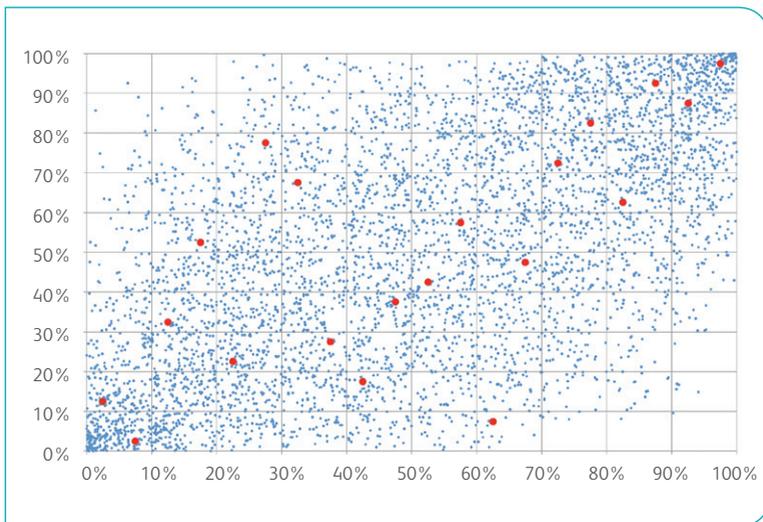
Appendix E

Validation tests for dependencies

Sense check

The “sniff test” is one of the most powerful validation tools for dependencies, as well as all other parts of the internal model. If it doesn’t stack up then there is likely to be a problem somewhere.

As an example, the following scatter plot is overlaid with the actual history as red dots. This supports the inclusion of a dependency since there seems to be a relationship between the two items, but also supports the use of a Student-t copula since there seems to be tail dependence in bottom left and top right of the graph. The use of historic data in this context can prove a powerful validation tool, which in this case, might invalidate the assumed relationship which shows minimal tail dependence since the bottom left and top right of the internal model outputs (the blue dots) just exhibit “noise”.



Independent assessment

The selection of dependencies requires challenge at various levels of the organisation. It's not only a second pair of eyes which is required but the combined knowledge of many throughout the organisation.

Following workshops with underwriters, actuaries and different business stakeholders, a common practice is to ask the chief actuary, chief underwriter and CRO to make an independent assessment of key dependency assumptions and to compare results and explore the reasons for differences in views.

“Being able to visualise dependencies makes them easier to understand and review. This applies to both the strength of the relationship and the level of tail dependence.”

Stress and scenario testing

Stress and scenario testing is an important part of the validation of dependencies because they are easy to understand and challenge. Their specific use in dependency validation could be along the lines of:

- Identify scenarios that include a number of dependent items. These could be based around individuals' past experiences (e.g. World-Trade Centre giving rise to large insurance losses across many classes), or using some of the previously discussed concepts such as causal analysis.
- Develop the financial impact of the scenarios in a workshop of business experts.
- Assess where the loss amount appears on the appropriate loss distribution curve from the internal model, and reverse-engineer the probability of a loss of that magnitude.
- Review and challenge this implied probability (both within the workshop, but also wider at Board review meetings). Does it make sense? How does it compare against any historic events?

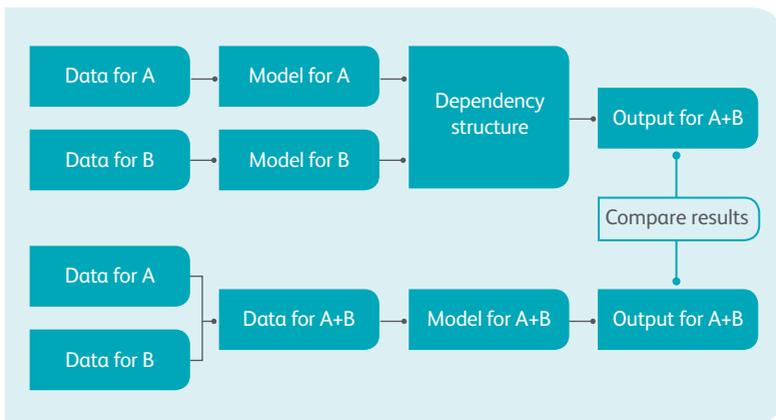
Another approach to using scenarios to validate the dependencies is to investigate the extreme simulations within the internal model (particularly around the 99.5 % ile for the purposes of Solvency II). What components of risk make up these simulations? What is present that you wouldn't have expected? What isn't present that you would have expected to be? Is there a consistent driver for these simulations, and if so, have you carried out enough validation around this key assumption?

“Common sense goes a long way when trying to validate dependencies because of the limited data. And the combined common sense of many should improve the outcome.”

Back-testing and aggregation

Dependencies do not manifest themselves in historical data very easily; however, back-testing can still be applied. For example, the historical performance of EL and PL business can be combined, and then compared against the overall internal model loss curve for the combined portfolio. If they are materially different, then have you got the underlying EL and PL curves appropriately calibrated? If they are, then the back-testing would suggest the dependency structure isn't appropriate.

Aggregation is slightly different, and involves carrying out combined modelling of items. The following schematic shows that this looks like when validating the dependency structure between two classes; A and B.



If the two different approaches give a very different answer, then you need to investigate why. It might be that you have modelled A or B poorly, or it might be that the dependency between them isn't well reflected, but either way, it would warrant the additional work to understand it better.

Benchmarking

Various consultancies and industry bodies publish benchmarks covering:

- Correlations between risk factors or lines of business
- Overall split of capital by risk type
- Ratios of capital to exposure measures

These can be helpful in assessing the level of dependency and aggregate capital within the model. There are also market studies on dependencies, such as the AON Benfield Insurance Risk Study (http://thoughtleadership.aonbenfield.com/Documents/20140912_ab_analytics_insurance_risk_study.pdf), and a number of regulatory models that show their dependency assumptions, such as the Solvency II standard formula and the Australian regulator's (APRA) Prescribed Capital amount (<http://www.apra.gov.au/CrossIndustry/Documents/141120-GPS-110.pdf>).

How well do your dependencies compare against the market? How do they compare against regulatory models? Can you explain why your dependencies should be different? What is the impact of running your internal model using benchmark or regulatory dependency structures?

Input and output tests

Checks should be made to confirm that the output dependencies look the same as those being put into internal models. Some software packages make changes to the dependency assumptions if they are fed with incorrect inputs (to use the technical term, negative semi-definite correlation matrices). A simple reconciliation here ensures that the model is doing what it says on the tin.

Tail dependency tests

Understanding the strength of the relationship in the tail is of key importance as this is a key determinant of the implied capital need. It is easier to try and validate day-to-day correlations than the more extreme tail correlations. However, there are a number of validation tools that can support tail dependency analysis.

The best tail dependence test is to visually represent dependencies and tail dependence, including scatter plots, contour maps and density maps such as the one on the opposite page. This shows the percentage of simulations that fall into different deciles of the joint distributions, and demonstrates that there is tail dependence between underwriting risk and reserve risk in this particular internal model. However, does it stack up against what you would expect? For example, if your underwriting is performing badly, then do you think it is materially more likely for your reserving to also experience deteriorations?

These density maps can also be linked through to independent challenge, with a variety of business experts asked questions such as "If underwriting was having a 1 in 10 bad year, then what chance do you think there is that reserving is also having a 1 in 10 bad year?" This can then be compared against the probabilities in the map.

Questioning along these lines can be extended to create "joint exceedance curves" that look at what happens to the correlation coefficient as you zoom into the top right-hand corner of a graph. The quicker they get to the zero then the less tail-dependence you have in your outputs. A related view would be to look at "conditional exceedance curves" curves to look how likely an extreme value of one variable is when the other one has poor experience.

If you do have sufficient data, then statistical validation tests can be employed using the density map. Actual data can be plotted alongside the expected data from the model (possibly using broader bandings than the 10 % shown in the example above). Chi-square tests can then be used to test how far the simulation differs from the historical data, and a confidence level for the modelled dependency being correct can be calculated.

Sensitivity testing

This is an important validation tool to help identify what is important. Flex some of the dependency assumptions and see what happens to the capital. From past experience, you might expect the correlation between credit risk and cat risk to be important, but once you sensitivity test these assumptions you might well find that it does not move the dial. However, something like the strength of the tail dependence parameter could easily swing the capital requirements by $\pm 10\%$. If this is the case, then have you spent enough time validating this important driver?

		Underwriting risk									
		0% – 10%	10% – 20%	20% – 30%	30% – 40%	40% – 50%	50% – 60%	60% – 70%	70% – 80%	80% – 90%	90% – 100%
Reserve risk	0% – 10%	3.0%	1.6%	0.9%	0.7%	0.5%	0.6%	0.4%	0.5%	0.4%	0.1%
	10% – 20%	2.1%	2.1%	1.3%	1.2%	0.9%	0.7%	0.5%	0.4%	0.2%	0.1%
	20% – 30%	1.1%	1.8%	1.6%	1.1%	1.1%	1.0%	0.8%	0.5%	0.4%	0.0%
	30% – 40%	1.0%	1.2%	1.5%	1.5%	1.5%	1.2%	1.0%	0.7%	0.6%	0.2%
	40% – 50%	0.6%	1.0%	1.3%	1.0%	1.4%	1.6%	1.2%	0.9%	0.6%	0.1%
	50% – 60%	0.5%	0.6%	0.9%	1.2%	1.3%	1.6%	1.5%	1.2%	1.0%	0.1%
	60% – 70%	0.4%	0.6%	0.7%	0.9%	1.1%	1.7%	1.6%	1.5%	1.0%	0.3%
	70% – 80%	0.4%	0.6%	0.7%	0.9%	0.7%	1.2%	1.8%	2.0%	1.8%	0.3%
	80% – 90%	0.2%	0.2%	0.4%	0.6%	0.4%	0.6%	1.1%	1.8%	3.1%	1.2%
	90% – 100%	0.1%	0.1%	0.0%	0.1%	0.1%	0.1%	0.2%	0.5%	1.3%	2.2%

Our project team

We would like to thank the following people and their employers for their work on this document:

Workstream leader:

Justin Skinner, Enterprise Risk Management Director, QBE European Operations

Consultancy support:

Russell Ward, Senior Consultant, Life Insurance, Milliman

Carl Gaffney, Consulting Actuary, Life Practice Milliman

Workstream members:

Jonathan Bilbul, Head of Capital Model Development at AIG

Edward Toman, Senior Actuary, ERM, Travelers

Our steering committee

The IMIF steering committee comprises:

José Morago, Aviva (Chair)

Tamsin Abbey, Deloitte

Raphael Borrel, Vitality

Simon Cartlidge, L&G

Alastair Clarkson, Standard Life

Robert Collinson, Towers Watson

Caroline Coombe, ORIC International

Karun Deep, RSA

Sebastien Delfaud, Bank of England (PRA)

Vishal Desai, Bank of England (PRA)

Steven Graham, Institute and Faculty of Actuaries

David Innes, RSA

Roger Jackson, KPMG

Tony Jeffery, Bank of England (PRA)

Rob Merry, HSBC

Niraj Shah, Ageas

Michael Sicsic, Aviva and ORIC International

Justin Skinner, QBE

Michael van Vuuren, EY

Barney Wanstall, PwC

Russell Ward, Milliman

Carolyn Williams, IRM

The Internal Model Industry Forum

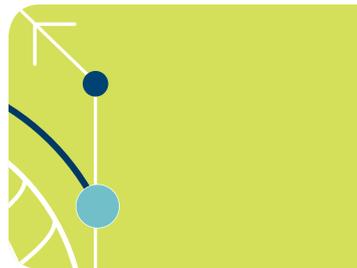
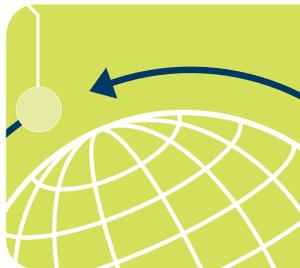
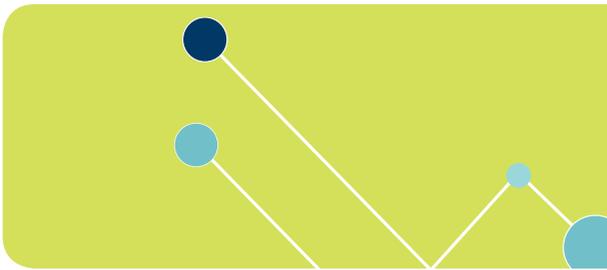
This document has been produced by the Internal Model Industry Forum (IMIF). The Institute of Risk Management (IRM) set up the IMIF in 2015 to address the key questions and challenges that insurers face in the use, understanding and validation of internal risk models. It is designed to work in a collaborative way to develop and share good practice to ensure that these models add value to the organisation and support regulatory compliance. IMIF now has over 300 members and we have run a series of Forum meetings to explore key issues. A number of workstreams are also undertaking research and we aim to publish the results along with other useful resources and guidance.

As the leading organisation promoting education and professional development in all aspects of risk management, IRM is pleased to be able to support this industry initiative to share good practice.

More information about the IMIF and its work can be found on the IRM website www.theirm.org

Who are the IRM?

This work has been supported by members of IRM, which has provided leadership and guidance to the emerging risk management profession for over 25 years. Through its training, qualifications and thought leadership work, which includes seminars, special interest and regional groups, IRM combines sound academic work with the practical experience of its members working across diverse organisations worldwide. IRM would like to thank everyone involved in the IMIF project.



IRM

T: +44(0) 20 7709 9808

E: enquiries@theirm.org

www.theirm.org

Institute of Risk Management

2nd Floor, Sackville House

143-149 Fenchurch Street

London

EC3M 6BN

United Kingdom

Our supporters

As a not-for-profit organisation, IRM is reliant on industry support to publish guidance like this. We would like particularly to thank the following organisations who have made this publication possible:



Barney Wanstall

barney.wanstall@uk.pwc.com

Bill Gasson

bill.gasson@uk.pwc.com

www.pwc.co.uk/



Russell Ward

russell.ward@milliman.com

uk.milliman.com/



Tamsin Abbey

tabbey@deloitte.co.uk

www2.deloitte.com/uk/en.html



Roger Jackson

roger.jackson@kpmg.co.uk

www.kpmg.com/



Michael van Vuuren

mvanvuuren@uk.ey.com

www.ey.com