Association of British Insurers

ABI RESEARCH PAPER 16, 2009

ANALYSING OPERATIONAL LOSSES IN INSURANCE

Evidence on the need for scaling from the ORIC® database

Report from ABI Research Department
By Mariano Selvaggi

FOREWORD

When the ABI established the Operational Risk Consortium Ltd (ORIC) in 2005, after several years of planning and design work, the intention was never to just collect loss data on operational risk. For although this task remains important, in operational risk, as with other key risks, the equally important task is to find analytical tools to process and manage risks. This paper *Analysing Operational Losses in Insurance,* published as an ABI Research Paper, is a significant contribution to that continuing work.

The complex nature and heterogeneous causes of operational risk – from IT failure to accounting and finance failures means this will always be a risk class where qualitative input will have a role, and the construction of numbers will require careful review. ORIC is committed to undertaking this work in parallel with collecting the data and using both to inform company practice and regulatory policy.

This work is more necessary than ever, as we consider the implementation of the Solvency II framework for insurers and the place of operational risk in both the standard approach and, where appropriate, in statistical models. A standard factor approach seems far from adequate, yet the most senior executives in the industry are calling for proportionality in Level 2 and 3 work – they do not want a framework that is so complex and cumbersome that makes risk management less rather than more effective. Operational risk professionals have the key role of developing methodological approaches that steer between these two extremes – this research paper should play a helpful role in supporting that work.

Peter Vipond
ORIC Chairman
Director of Financial Regulation & Taxation
Association of British Insurers

EXECUTIVE SUMMARY

Historical loss data are essential for the effective measurement and management of operational risks. Operational losses in insurance do not happen very often, though, and take time to crystallise. Internal losses are thus unlikely to provide a full picture of the spectrum of risks faced by the firm. They are a biased sample of the universe of potential losses because they reflect idiosyncratic features such as the firm's business, culture and control environment. Information on losses experienced by other firms can fill important gaps in this knowledge, but to deliver meaningful results it is vital these loss events are "comparable" to the losses the firm might experience.

This research studies robust methodologies for scaling the size and number of external losses to make them equivalent to a firm's internal loss events. Adjusting for potential scaling biases is important when external and internal losses are merged for operational risk management and economic capital calculations.

We use operational loss event in the ORIC database to provide real-world applications of the methodologies discussed. It is the first time we have used our data in this way. We set this against data on the size of the insurer where the loss occurred and additional scaling factors controlling for business lines and loss event types. The purpose of our research is not to provide final answers, but to illustrate our empirical approach and uncover early trends in operational loss data from insurance business.

Among the main conclusions of our empirical analysis are the following:

- The size of the insurer is strongly associated with the severity and number of its operational losses. More specifically, loss amounts are positively correlated with the number of full-time employees whereas the number of loss events in a given quarter is more sensitive to premium income.

- Increasing the number of full-time employees by 1% results in an increase of about 0.8% in the predicted loss amount, holding all other variables constant. The standardised loss amount per event, which according to our estimates is around £27,000, can increase to nearly £300,000 in firms with high number of full-time employees.

- For a standard deviation rise in premium income, roughly £3.7bn, the projected number of operational losses per quarter increases by 24 per cent, holding other variables constant.

- Insurers at the lower end of the distribution of premium income are predicted to experience more than nine losses only 4 out of 100 times, whereas insurers in the upper end of the distribution would experience more than nine losses 12 out of 100 times.

- The business functions Customer Service/Policy Administration and Claims are strongly associated with smaller losses, other things being equal. Business functions can often predict the variability in observed operational losses better than business lines.

- While the goodness-of-fit of our scaling models is often higher than the results reported by similar studies in the banking sector, a great deal of the variability in observed losses remains unexplained by our models.

- For scaling the frequency of operational losses, the negative binomial regression model was preferred to the commonly used Poisson regression model.

To ensure the robustness of these findings, we fitted additional scaling models using different sample sizes and econometric techniques (e.g. quantile regression). The final results did not change much.

ACKNOWLEDGEMENTS

I would like to thank my ABI colleagues for helpful suggestions and comments on early drafts of this paper, and Natasha Panjwani from ORIC for expert research assistance.

I am also extremely grateful to all ORIC members and insurers who helped throughout the preparation of this research, including:

- Aegon
- Allianz
- AXA
- Axeria Prévoyance
- Bupa
- Friends Provident
- HBOS
- Hiscox
- IAG
- Just Retirement
- Legal & General
- LV=
- NFU Mutual
- Pearl Group
- Prudential
- RBS Insurance
- Royal London
- RSA
- Scottish Widows
- Skandia
- Standard Life
- Swiss Re
- Travelers
- Unum

CONTENTS

LIST OF TABLES

LIST OF FIGURES

1.0 INTRODUCTION

Operational risk refers to the risk of incident or financial loss arising from inadequate or failed internal processes, or from personnel and systems, or from external events. Sound operational risk management (ORM) is increasingly important in the insurance sector, particularly since the introduction of the Individual Capital Adequacy Standards regime in the UK and the EU Solvency II Directive.[1] Operational risk is one the key risk modules in the Solvency II framework, for example, where the capital charge attached to this risk class could be up to 30% of the firm's capital requirement.[2]

Reliable historical loss data is essential for ORM and modelling of risk exposures. Yet gathering good-quality information on operational losses in the insurance industry is a challenge – arguably a bigger challenge than in banking activities. External loss events occurring in comparable firms are often mixed with internal losses to assist firms' modelling, but this approach raises methodological issues because the size and number of operational losses depend of idiosyncratic features such as the insurer's size, risk profile and control environment. Operational risk can be highly influenced by the firm. Failing to adjust for these factors when mixing internal and external losses may induce biases in the analysis, which lead to inaccurate estimates and misleading conclusions.

Using unique data from the Operational Risk Consortium (ORIC) database, we consider robust methodologies to "scale" the size and number of operational losses reported by members of the consortium to make them more comparable to a firm's internal losses. We identify statistically robust exposure metrics that work well to calibrate external losses and thereby ameliorate the potential effect of scaling biases.[3] As a result, external loss data can be combined with internal data in a rigorous systematic way.

For example, we find that:

- The size of the insurer, measured by premium income and number of full-time employees, is a statistically important determinant of the size and number of its operational losses, other things being equal. Loss amounts are more sensitive to headcount, while the predicted frequency of operational losses is more sensitive to gross premiums.

- Our estimates suggest that predicted losses per individual loss event, which for a hypothetical "standard" firm are around £27,000, could increase to £300,000 per loss event in firms with high number of full-time employees, everything else equal. We also find that for a standard deviation increase in premium income,

[1] Solvency II includes legal risks in its definition of operational risks, but excludes risks arising from strategic decisions, as well as reputational risks.

[2] The other three risk modules are underwriting, market and credit risk. See, for example, Doff (2007) and Sandström (2009) for a thorough discussion of operational risk within the Solvency II Directive.

[3] Combining loss data from several firms to build a predictive model for the operational risk of a single firm raises "scaling" problems, because firms are of different size and face asymmetric risk exposures.

roughly £3.7bn, the predicted number of loss events per quarter could increase by 24 per cent, holding other variables constant.

- The exponential distribution is the best fit to observed operational loss amounts, but the lognormal distribution is preferred when one considers “scaled” losses. This result has important implications for the calculation of capital buffers linked to operational risk using internal models.

- Firms at the lower end of the size distribution by premium income are predicted to have more than nine losses a quarter only 4 out of 100 times, whereas firms in the upper end of the size distribution would experience more than nine losses a quarter 12 out of 100 times.

- The negative binomial regression model represents a better fit to the observed number of operational losses than the Poisson regression model.

1.1 Severity and frequency of operational losses

Following existing literature, we examine separately probability distributions of the loss frequency and loss severity. This analysis involves discrete and continuous statistical distributions, respectively, which are assumed to be independent of each other. Having a good understanding of both numbers and amounts of losses is important not only for modelling purposes, but also to identify effective management and mitigating actions.

In practice, frequency and severity distributions are then “combined” using stochastic simulations and correlation assumptions (the so-called *convolution* process) to develop models of aggregate losses, which are used to calculate economic or regulatory capital for operational risk. Here we omit this last step because some of our empirical results, while statistically significant, are considered preliminary and susceptible to change. A risk measure such as value-at-risk for a (hypothetical) standardised firm may be taken out of context and induce unintended misinterpretations. Therefore we focus on scaling methodologies here, leaving the issue of convolution for future research.

1.2 Structure of the paper

The rest of the paper proceeds as follows. Section 2 describes general features of ORIC and provides a high-level introduction to ORIC’s database. This section also describes the sample of ORIC loss events used in our applied analysis. Section 3 summarises the main sources of operational loss data and discusses the challenges that arise when information on external loss events is mixed with internal losses. Sections 4 and 5 represent the core of our empirical analysis. They analyse scaling and curve fitting for the severity and frequency of operational losses, respectively. The appendices contain additional technical details.

2.0 THE OPERATIONAL RISK CONSORTIUM (ORIC)

The Operational Risk Consortium (ORIC) is the consortium for the insurance and asset management industry that collects, standardises and reports operational risk loss data. High-quality historical information on loss events is very important for the successful measurement and management of operational risk. ORIC embodies best practice and a quality-controlled database to support firms' solvency capital calculation and risk management activities, with the ultimate goal of making more informed business decisions.

ORIC was created in 2005 by the Association of British Insurers (ABI) together with 16 core insurers to help the industry improve its understanding of operational risk. ORIC remains a not-for-profit organisation, and its current members are drawn from both life and non-life business. Figure 1 shows the distribution of gross loss amounts, by business function and for the life and non-life business units.

Figure 1 Distribution of operational loss amounts as at 2009:Q1, by business function and business unit

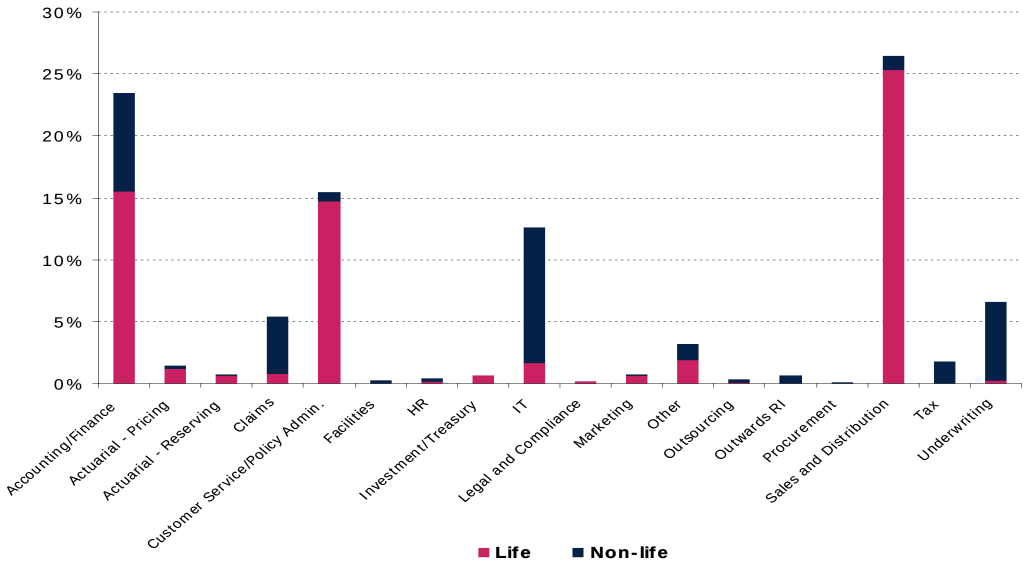

Note: We excluded from this analysis the event category "Mis-selling (Endowments)".

Source: ABI Research based on ORIC database.

While ORIC is an independent entity, the ABI keeps its role of managing company as it is uniquely placed to provide impartial, centralised and dedicated resources for the insurance and investment industry.

2.1 How ORIC supports firms' activities

ORIC contributes to its members' activities along several dimensions that involve both quantitative and qualitative aspects of their business:

- **Quantitative analysis**. A key feature of ORIC is its operational loss event data, which in turn translates into
  1. Timely reporting;
  2. Sound risk analytics/intelligence; and
  3. Robust operational risk modelling (loss distribution and scenario approach).

- **Think tank**. ORIC also represents a unique forum for risk managers throughout the industry to get together and analyse trends in their business. This supports the identification of shared standards and sound practices for the measurement and management of operational risk in insurance.

- **Workshops**. ORIC organises regular workshops on relevant "hot topics" facing the industry (e.g., scenario analysis). These show-and-share sessions involve a great deal of cross-learning and are ideal environments for the identification and dissemination of best practices.

2.2 ORIC and the regulatory landscape

ORIC represents an industry-led response to modern regimes for prudential regulation. Particularly the UK Financial Services Authority (FSA) and Solvency II, the EU Directive recently approved by the European Parliament, both consider operational risk exposure when defining solvency requirements for insurers.

ORIC participates actively in regulatory discussions to represent its members and also promote good risk-management practices within the industry. ORIC sits on the FSA's Operational Risk Expert Group for insurance, which discusses topics linked to the local supervision of insurers and implementation of Solvency II. At the European level, ORIC also maintains dialogue with the Committee of European Insurance and Occupational Pensions Supervisors (CEIOPS) to ensure its loss database fulfils the statistical quality, suitability and completeness required by CEIOPS to assist firms in their internal model approval process under Solvency II.

A number of recent discussion papers and feedback reports published by the FSA and CEIOPS identify ORIC as an important industry-wide classification standard and source of external information on operational events to improve the completeness of insurers' internal loss data.[4]

Estimating appropriate capital requirements for operational risk is a challenge because insurers' internal data on operational loss amounts and near misses are often limited and biased. A reliable and complete source of external loss information assists the prudent assessment of capital charges for operational risk. Consortium-based loss data also provides an adequate benchmark against which a firm's own loss experience and risk management framework can be compared.

[4] See, for example, FSA (2008a, 2008b) and CEIOPS (2008).

2.3 Overview of the ORIC database

The ORIC database comprises operational risk loss data – information on loss amounts and incidents due to failures in people, processes, systems or external events. A loss event represents an incident that causes a (direct or indirect) financial loss to the firm. The lower threshold for operational loss amounts in the ORIC database is £10,000.

ORIC contains in-depth narratives of the circumstances surrounding the incident and a hierarchical classification system for event types and root causes of the loss. ORIC can accommodate global submissions from different geographical areas and currencies. It also captures "near misses", as they contain valuable information for the qualitative management of operational risk. These are cases where operational controls failed but financial losses did not materialise. Near misses can be quantifiable or unquantifiable.

As at 2009:Q1, there were nearly 2,000 incidents in the ORIC database representing a gross loss amount of almost £900m. These losses refer to insurance business only, but ORIC intends to also roll out its platform for asset management shortly.

2.3.1 Categorisation of losses

ORIC has set standards for event type classification. The Level 1 and 2 categories are consistent with the new Basel accord. ORIC has also developed a Level 3, which breaks the Level 2 classification into over 70 further categories to improve the granularity of the data. Figures 2 and 3 show "risk maps" that combine the frequency and severity of operational losses for each Level 1 event type. (We excluded the event type misselling (endowments) because loss events are captured differently and it skews the analysis.)

Figure 2 Risk map of gross losses, by Level 1 category

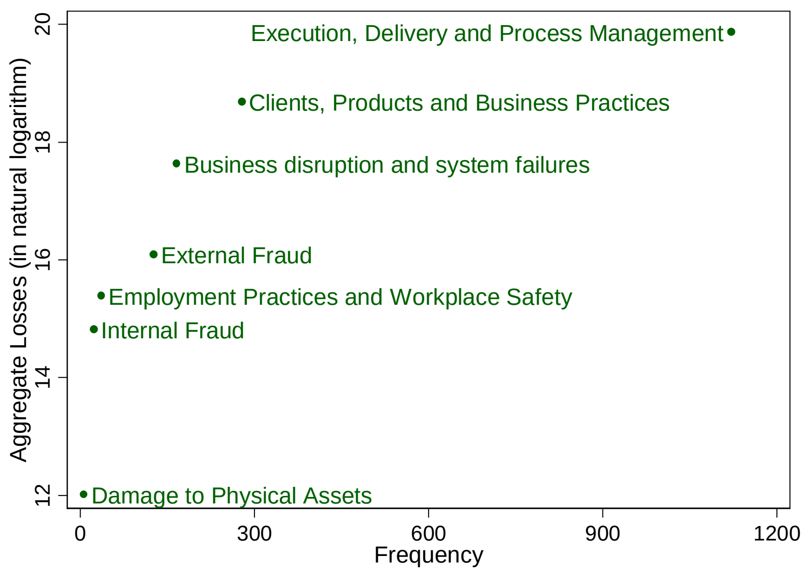

Note: We excluded from this analysis losses before 2000 and the event category "Mis-selling (Endowments)".

Source: ABI Research based on ORIC database.

Figure 3 Risk map of gross loss per case, by Level 1 category

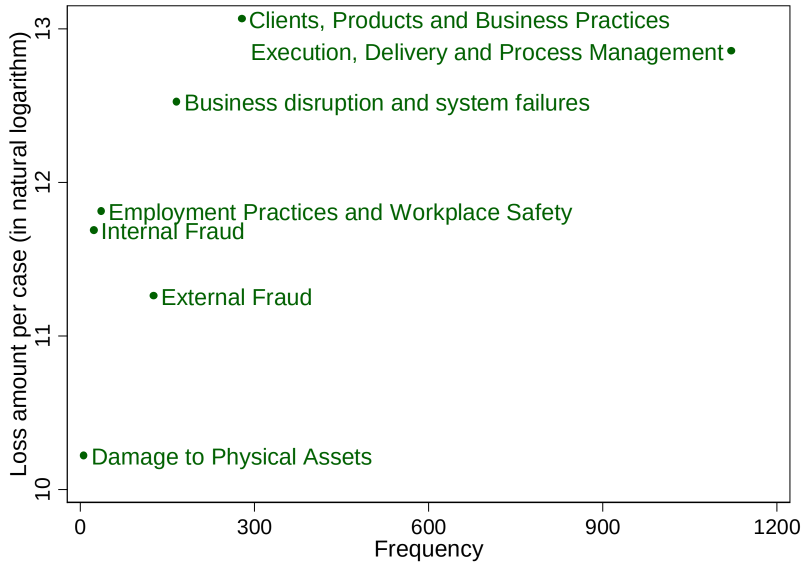

Note: We excluded from this analysis losses before 2000 and the event category "Mis-selling (Endowments)".
Source: ABI Research based on ORIC database.

Figure 4 shows "heat maps" for the severity and frequency of gross operational losses. They provide an idea of the concentration of loss events across business functions and risk event types. As can be observed, the impact is concentrated in the loss event type *Execution, Delivery and Process Management*, and in the business function *Sales and Distribution*. Customer-related and advisory activities in turn concentrate a significant portion of the number of losses in the ORIC database.

Figure 4 Heat maps for severity and frequency of ORIC losses

4A – Distribution of the size of operational losses

Level 1 category / Business Function	External Fraud	Clients, Products and Business Practices	Execution, Delivery and Process Management	Employment Practices and Workplace Safety	Business disruption and system failures	Internal Fraud	Total
Sales and Distribution	0.4 %	18.9 %	6.8 %		0.3 %		26.4 %
Facilities			0.1 %				0.2 %
Customer Service/Policy	0.7 %	1.2 %	13.2 %		0.3 %		15.5 %
Investment/Treasury			0.7 %				0.7 %
Accounting/Finance			23.4 %				23.5 %
Other		0.1 %	2.3 %	0.4 %	0.1 %	0.1 %	3.2 %
IT		2.7 %	3.8 %		6.0 %		12.6 %
Legal and Compliance			0.1 %				0.1 %
Tax			1.7 %				1.7 %
Claims	0.7 %	0.2 %	4.0 %		0.3 %	0.2 %	5.4 %
Outsourcing			0.3 %				0.3 %
Underwriting		0.1 %	6.3 %				6.6 %
Marketing		0.5 %	0.2 %				0.7 %
Actuarial- Reserving		0.1 %	0.5 %				0.7 %
HR			0.2 %	0.2 %			0.4 %
Actuarial- Pricing			1.2 %		0.2 %		1.4 %
Procurement							
Outwards RI			0.6 %				0.6 %
Total	1.8 %	24.0 %	65.5 %	0.8 %	7.4 %	0.5 %	100.0 %

Greater than 5 %
Between 1 % and 5 %

4B – Distribution of the size and number of operational losses

Event Category Level 2	Event Category Level 3	Severity	Frequency
Advisory Activities	Mis-selling (other)	13%	9%
Transaction Capture, Execution and Maintenance	Accounting Error	12%	2%
	Inadequate process documentation	8%	3%
	Transaction System Error	8%	6%
	Management Information Error	7%	1%
	Data Entry Errors	7%	5%
	Management failure	5%	2%
	Customer Service Failure	4%	16%
Suitability, Disclosure and Fiduciary	Customer Complaints	6%	4%
Systems	Software	6%	3%
Vendors and Suppliers	Vendor Delivery failure	3%	2%
Product Flaws	Product Design	3%	1%
Customer or Client Account Management	Incorrect Payment to Customer/Client	2%	9%
	Payment to Incorrect Customer/Client	1%	4%
Theft and Fraud	Fraudulent Claims	1%	4%

Greater than 10%

Between 5% and 10%

Note: We excluded from this analysis the event type category "Mis-selling (Endowments)". Figure 4B focuses on a subsample of Level 2 and Level 3 categories, which includes the most relevant event types.

Source: ABI Research based on ORIC database.

2.3.2 Our ORIC sample

The empirical analysis discussed in Sections 4 and 5 is based on a sample of losses in the ORIC database. Particularly we focus on operational loss events reported by 18 insurers that occurred between 2005 and 2008. Our sample comprises 1,388 events for a total loss amount of about £380m. The firms' combined annual premium income is, on average, £60bn. The following table shows some basic statistics for all the losses and for those related to the life and non-life business units.

Table 1 Descriptive statistics for ORIC sample

	No of events	Percentage	Average size of loss (£000)	Aggregate size of losses (£m)
Total losses	1,388	100%	273	379
Life losses	1,056	76%	259	274
Other losses	332	24%	316	105

Note: Other losses refer to the General and Group business units.

Source: ABI Research based on ORIC database.

Figure 5 shows the distribution of loss amounts across the 18 firms in our sample. Each box contains observations between the 25th - 75th percentiles of the distribution whereas the dividing line across each box corresponds to the median. The dots in turn represent outsiders or very extreme values. The biggest loss amount per event in our sample is £41m.

Figure 5 Distribution of loss amounts, by insurer

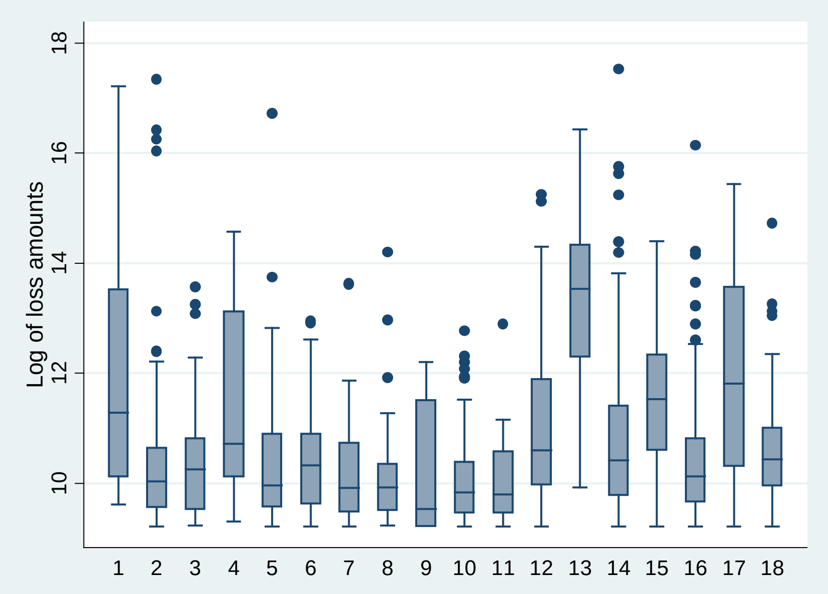

Source: ABI Research based on ORIC database.

3.0 SOURCES AND CHALLENGES OF LOSS EVENT DATA

The lack of data on actual losses is one of the main stumbling blocks when modelling operational risk exposures. This issue has already been faced in the banking sector but is likely to have a bigger impact on the insurance industry because operations involve fewer transactions and less trading, which are important drivers of operational failures.

A natural way to overcome the issue of lack of data about low-frequency-high-impact events is to combine data from different sources. Specifically, an external database of losses that took place in comparable firms could be merged with internal loss data for modelling purposes. However this data aggregation exercise poses challenges because companies have different business profiles, sizes and controls that are likely to affect the size and frequency of their operational losses. Failing to control for these factors when mixing loss data may lead to flawed estimates and misleading conclusions.

In this section, we first present the most important sources of information on external loss events – publicly available losses, proprietary and consortium-based databases. Then we discuss the key challenges raised by each data source, and argue that consortium-based loss data are the best source of external data and can be adjusted to help firms overcome scaling issues that naturally arise when internal and external operational loss events are combined.

3.1 Sources of operational loss data

In analysing effective ways to combine internal and external loss data, it proves useful to start by identifying the main sources of operational risk loss data:

1. **Internal loss data**. An internal operational loss database is a key input into any effective risk management programme. However internal data on operational losses are often limited and potentially biased, not only due to internal collection problems but also because large operational losses in insurance do not happen very often and take time to crystallise. Internal historical loss events reflect the specific experience and history of the organisation, and are therefore unlikely to provide a complete picture of the universe of major operational risks the firm faces. Sources of information on loss events that occurred in comparable firms can sometimes prove helpful to fill the gaps in this key knowledge.

2. **Publicly available losses.** These databases contain operational risk losses that, because of their magnitude or high profile, filter through to the public domain. Information about incidents is obtained from public sources such as (printed and electronic) news reports, specialised publications, regulator's news releases and court filings, so these loss events tend to be linked to low-frequency-high-impact operational failures. IT vendors such as Fitch's OpVantage, Open Pages and SAS offer this kind of external loss database. Willis, the professional services firm and

insurance broker, also created a database of public operational risk loss events from the financial services industry.

3. **Commercial database using proprietary loss data.** These databases contain information on operational risk loss events experienced by financial institutions that are not necessarily in the public domain. For example, some of the incidents may relate to insurable losses or the data may have been extracted from insurance claims files, irrespective of whether a claim was made by the company or not. Aon, the risk management and insurance brokerage firm, manages a database that combines proprietary loss data not available elsewhere with loss information sourced from the public domain.

4. **Consortium-based loss data.** These databases comprise loss events reported to a consortium by its members, who in return get access to anonymised, pooled industry data on operational loss events and near miss incidents. The members of the industry consortium pay a subscription fee and voluntarily commit to feed the database with their individual internal loss events so long as confidentiality of key information is protected. ORIC (for insurers), and BBA's Gold and ORX (for banks) fall into this category.

The threshold above which public losses are recorded is likely to be much higher than the reporting threshold for consortium data. For example, ORIC has a lower threshold of £10,000 while some public loss databases have lower thresholds of at least £1m. It is also worth mentioning that most of the existing databases have loss amounts and incidents linked to banks. Even databases with operational risk incidents in financial institutions are severely skewed towards banking. ORIC is the only consortium-based operational loss database dedicated to the insurance industry.

3.2 Challenges raised by external loss databases

The use of external loss data for modelling purposes poses some real challenges. Since the severity and frequency of operational losses are likely to depend on firm-specific characteristics such as its risk profile, size or control environment, it is not advisable to merge indiscriminately internal data and losses experienced by other firms. Certain preliminary calibration is needed to ensure the approach is applied rigorously.

The source of information on external losses may also have implications for the quality of the analysis. The way in which loss databases are constructed and operational risk events categorised is also important, since inconsistencies may lead to potential biases that need to be adjusted for when peers' losses are mixed with internal data.

The number and magnitude of potential problems vary according to the source of data being used. Some external loss events are more susceptible to biases than others. This is a function of the underlying collection and classification processes. Table 2 describes some of the most common biases and the situations in which they are likely to arise.

Table 2 Biases affecting external op risk loss databases

Type of bias	Description of the problem	Main impact on
Controls	The number and severity of operational losses are likely to depend on each firm's risk control environment and risk management process. To improve the reliability of final estimates, it is important these differences are taken into account when external losses are mixed with internal data to model an institution's exposure to operational risks.	All databases
Scaling	This captures the notion that the size of an operational loss is often correlated to the size of the firm suffering the loss. Indeed the frequency of operational losses experienced by a firm may also be sensitive to the size of the firm. Some form of scaling of peers' loss events is therefore needed in order to make external data properly comparable to the internal losses of the company under consideration.	All databases
Selection	This refers to the fact that certain operational losses may be more or less likely to be publicly reported because of their severity or importance. Size and type of losses can be drivers of the underlying selection mechanism. The losses in the database are those that cannot escape from public scrutiny and are thus 'self-selected', which gives rise to the random (unobserved) truncation problem - the process by which public losses are picked up is not known.[5]	Public loss databases, and to a lesser extent proprietary databases
Collection/ Reporting	This form of selection bias arises when the losses in the database are not a random sample of the population of all operational losses, but instead a biased sample containing a disproportionate large number of losses fulfilling specific reporting or collection criteria. The size or type of loss can be a factor affecting the selection mechanism. Problems arise when one relies on these databases to make out-of-sample inferences for the wider population. Collection biases may also arise when external data points from different sources are pooled together due to possible different truncation points and thresholds in the databases.	Public and proprietary loss databases
Classification	Operational loss events may not be consistently defined or classified within the database. For example, similar loss events may be misclassified under different loss types.	Public loss databases

5 De Fontnouvelle et al. (2003) is an early paper discussing econometric techniques to tackle selection biases due to the use of publicly available operational loss databases.

Databases sourced from the public domain are likely to introduce more potential biases than consortium-based databases, say. Very often this simply reflects features of the collection mechanism being applied or the quality of the validation and classification of operational loss events submitted to the consortium.

Databases sourced from the public domain can sometimes hamper the correction of biases because not all the required information is readily available within the database. Exposure metrics may be difficult to identify. One advantage of consortium-based loss data is that close dialogue with member firms facilitates the identification of business areas originating the loss event and improves the precision of scaling techniques.

4.0 SEVERITY ANALYSIS

Information on low-frequency-high-impact losses is extremely important for estimating prudent capital buffers for operational risks. Yet internal loss databases are unlikely to provide a complete picture of the spectrum of risks because, almost by definition, they are a small and potentially biased sample of the universe of major operational risks the firm faces. While external loss events are helpful to improve data coverage, they need to be handled with care because external losses are not directly analogous of the firm's internal losses. Some scaling work that controls for characteristics of the external loss events is needed to make external and internal losses more comparable.

In this section we consider approaches to scaling the severity, or size, of operational losses. These loss events are defined as operational incidents where (direct or indirect) monetary losses occurred. We also examine two applications that show how to merge internal loss amounts with properly "scaled" external losses sourced from ORIC's consortium-based database.

Section 4.1 sets out our analytical approach, Section 4.2 discusses a methodology for scaling consortium-based loss amounts and lists our main findings, while Section 4.3 tackles the issue of curve fitting. In summary, we find that:

- The size of the insurer is positively correlated with the size of operational losses. Both premium income and number of full-time employees have positive impacts on the size of the loss, although only the latter effect is consistently statistically significant.

- The severity of operational financial losses is more sensitive to headcounts than to premiums.

- Customer- and claims-related business functions are negatively associated with the size of operational losses, other things being equal. Advisory activities are in turn associated with bigger losses after controlling for other scaling factors.

- When "scaled" external losses are combined with internal losses for curve-fitting purposes, the severity model with the best statistical fit is significantly different from the best one when only internal and "raw" external losses are used. Using scaled external loss data for severity distributions has important implications for regulatory capital calculations for operational risk.

4.1 Analytical approach

The so-called severity models studied in this section are built on continuous probability distributions, where the variable of interest is the operational loss amount. Losses can be of any size but in reality they are normally truncated at a lower threshold.

It proves helpful to partition our approach to severity models into two areas:

1. Scaling the severity of external loss events to make them comparable to the own losses of the firm under consideration; and
2. Fitting probability functions to internal and properly scaled external loss amounts to select the most appropriate statistical distribution for the size of losses.

4.2 Scaling model for the size of losses

This section looks at rigorous ways to make losses sourced from a consortium-based database more "comparable" to a firm's internal losses. Given that an operational loss occurred, we assume that its size depends on general as well as firm-specific factors such as the firm's size and control environment. General, or common, factors affect all firms similarly. Firm-specific, or idiosyncratic, factors may in turn vary across firms and loss events. It is precisely the failure to adjust for idiosyncratic elements through scaling of external losses that leads to wrong estimates and flawed conclusions when internal and external loss data are merged.

To scale external losses up or down, as appropriate, we use an econometric approach based on regression techniques. The details of this analysis are discussed in Appendix A.1. Simply put, we assume the size of operational losses depends on *general* and *idiosyncratic* factors. The former capture the context in which firms operate, such as the prevailing macroeconomic conditions or regulatory environment, whereas the latter capture characteristics of the firm or loss event. Idiosyncratic elements are also called "exposure metrics" in the literature, and are essential inputs to the scaling process.

In our model, the general component determines the size of a baseline operational loss value that is then scaled up or down by the value of the idiosyncratic component. We consider five main types of exposure metrics (or regressors):

- **Size of the insurer**. Measured by average annual premium income ("GWP") and average number of full-time employees ("FTE") over the period considered. We consider single and regular premiums, where appropriate, while full-time employees exclude contractors and outsourcers.

- **Firm-specific effects**. Indicator variables that identify the 5 insurers with more than 100 individual loss events in our ORIC sample. The purpose is to capture unobservable idiosyncratic characteristics that may be relevant to explain the variability in observed loss amounts.

- **Recovery**. Indicator variable ("Recovery") that takes value 1 if the event had a positive recovery amount, and 0 otherwise.

- **Business line**. Indicator variable ("Life") that takes value 1 if the loss occurred in the Life business unit, and 0 otherwise.

- **Business function**. Indicator variables identifying the most recurrent business functions in which the losses occurred: Customer Service/Policy Administration ("CS"), Sales and Distribution ("S&D"), and Claims ("Claims").

- **Loss event types**. Indicator variables for the following (Level 2) categories: Advisory Activities ("Advisory"), Customer Account Management ("Customer"), Systems ("Systems"), and Transaction Capture, Execution and Maintenance ("Transactions").

Our scaling mechanism for calibrating operational risk losses relies on a specific scaling law that is assumed to govern all firms' losses. Specifically, we scale loss amounts with the above variables to improve the comparison of operational risk losses from different insurers. The regression techniques discussed in Appendix A.1 help us asses how well our modelling assumptions fit the observed data.

4.2.1 Findings of the scaling model for the size of losses

The principal findings of this part of our empirical analysis are set out below.[6]

Box 4.1 Results for severity scaling

- The size of the insurer is strongly correlated with the size of the loss. We find that GWP and FTE are generally positively correlated with the size of the loss, although only full-time employees are consistently statistically significant.

- The business functions Customer Service/Policy Administration and Claims are strongly associated with smaller loss, other things being equal. These business functions are more statistically significant than the business line to explain the size of operational losses.

- The (Level 2) event types Advisory and Customer are strongly positively and negatively correlated with the size of the loss, respectively.

- Reassuringly, the estimated models are highly significant and the proportion of the variation in observed losses explained by the models (R-squared measure) is generally higher than those reported by similar studies in the banking industry.

Robustness tests

- To check the robustness of our findings we estimated two additional regressions, one using net losses (gross loss amount minus recovery amount) as dependent

[6] Detailed econometric results can be found in Appendix A1.

variable and another explicitly accounting for the fact that ORIC losses are truncated (from below) at the level of £10,000. The signs and significance levels of the coefficients did not change much, although in the truncated regression the individual effects of the explanatory variables identified above generally increased in magnitude.

- We also fitted quantile regression models, and concluded that the impact of the insurer's size (measured by GWP and FTE) on the severity of operational losses does not depend on the portion of the distribution examined. Specifically, only FTE has a significant positive impact on the size of operational losses across the whole distribution of losses, although its impact is diminished at lower quantiles.

- The event category Advisory is negatively correlated with loss amounts at lower quartiles of the severity loss distribution.

The results for GWP and FTE are robust although they should be interpreted with care because these two regressors are positively correlated, which makes it is difficult to tease out their individual effects on the size of losses (see Figure 6). Yet, taking both effects together we may conclude that the larger the insurer is, the more severe its operational losses tend to be.

Figure 6 Relationship between premium income (GWP) and full- time employees (FTE)

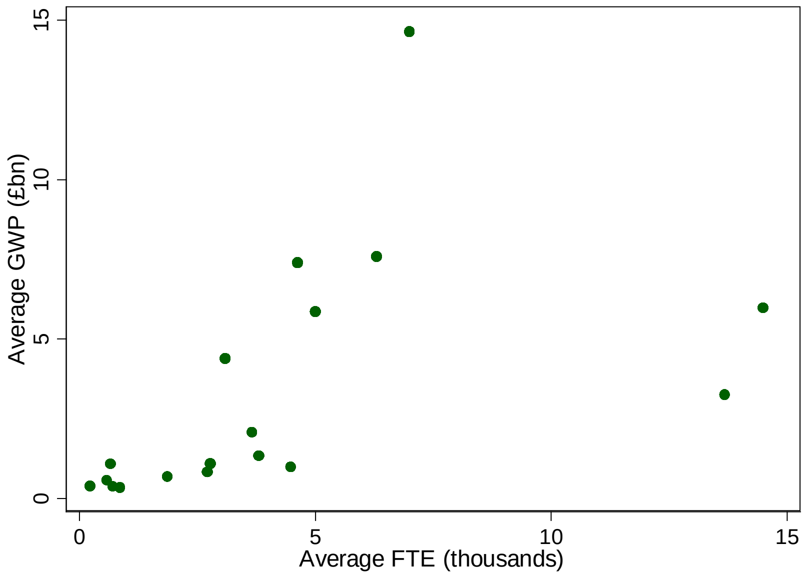

Source: ABI Research based on ORIC database.

Figure 7 shows the main results of our econometric analysis. Specifically it depicts the impact an increase in the respective explanatory variable has on the predicted size of an operational loss, when other things are held constant.

For the continuous regressors premium income and number of full-time employees, the results are given in the form of elasticity: the percentage change in the size of the loss for a 1 per cent change in the explanatory variable. For example, increasing by 1% the number of full-time employees leads to an increase of more than 0.7% in the predicted loss amount, holding all other variables constant. In the case of indicator variables, the results are presented in the form of expected number of standard deviation changes in the log-transformed loss amount. To illustrate, having a recovery amount results in an increase in the dependent variable of 0.34 standard deviations (or 1.34 log-transformed GBP amount), holding all other variables constant, whereas the customer services function results in an expected drop in the dependent variable of 0.36 standard deviations.

Figure 7 Impact of explanatory variables on the size of losses

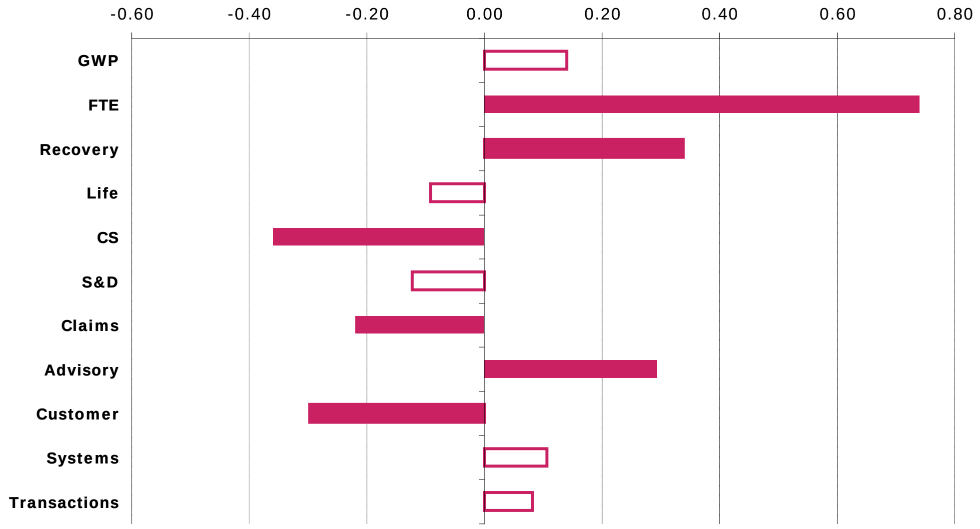

Note: Solid bars identify variables that are statistically significant at least at the 10% level.

In our econometric model, the value of a “standardised” loss amount is then re-scaled using the exposure metrics to predict operational losses that take into account specific features of the insurer under consideration, such its size. To illustrate, Figure 8 depicts the impact of premium income and number of full-time employees on the scaled value of this loss. As can be seen, predicted loss amounts per event are much more sensitive to the value of full-time employees than to premium income (indeed the effect of the latter is not statistically significant). The standardised loss amount, which according to our estimates is around £27,000, can reach nearly £300,000 per event for relatively high values of the full-time employees.

Figure 8 Estimated effect of GWP and FTE on re-scaled loss amount

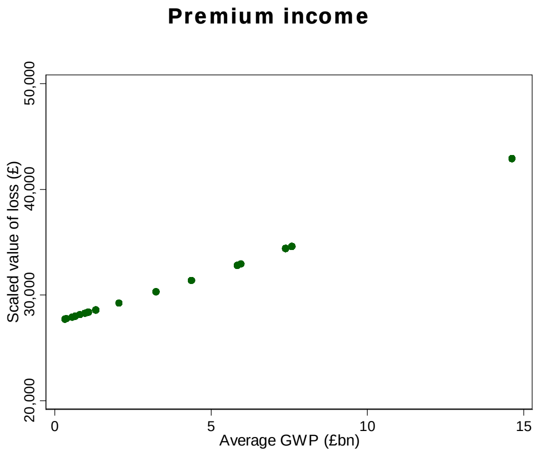

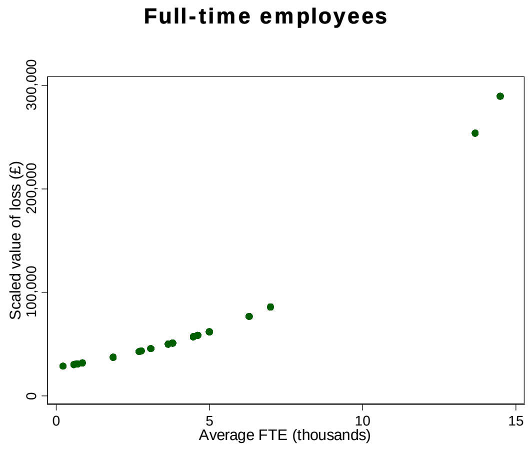

4.3 Curve fitting for the size of losses

In this section we concentrate on continuous probability functions in order to identify distributions that offer a good fit to the observed pattern of the size of losses. This is important because, in practice, the selected severity distribution feeds into stochastic simulations for aggregate loss amounts that then determine economic and regulatory capital buffers for operational risks.

Using the econometric estimates discussed above, we also illustrate how scaled losses originating from other insurers can be mixed with internal information on loss events to improve the accuracy of the curve-fitting exercise. Since we are only interested in operational losses, we restrict ourselves to statistical distributions that are not defined for negative values – positive loss amounts would entail gains.

Within this class of distributions, we focus on those that have been most widely used in the operational risk literature:

- The beta (generalised) distribution;
- The exponential distribution;
- The log-normal distribution;
- The Pareto distribution.

The functional forms of the above distributions are in Appendix A.2. A common feature worth mentioning here is their asymmetry or, more precisely, the fact they are heavily skewed; since small losses are more frequent than extreme losses, the mean loss per event is greater than the median loss amount. This property is observed very often in the distribution of real-world operational losses, which tend to be highly (positively) skewed.

4.3.1 Using scaled data to estimate the distribution of loss amounts

Figure 9 shows the empirical distribution of the size of actual and standardised losses, where the latter predictions are based on the estimates of our main econometric model (see Appendix A.1).

Figure 9 Distribution of actual and standard losses per loss event (in logs)

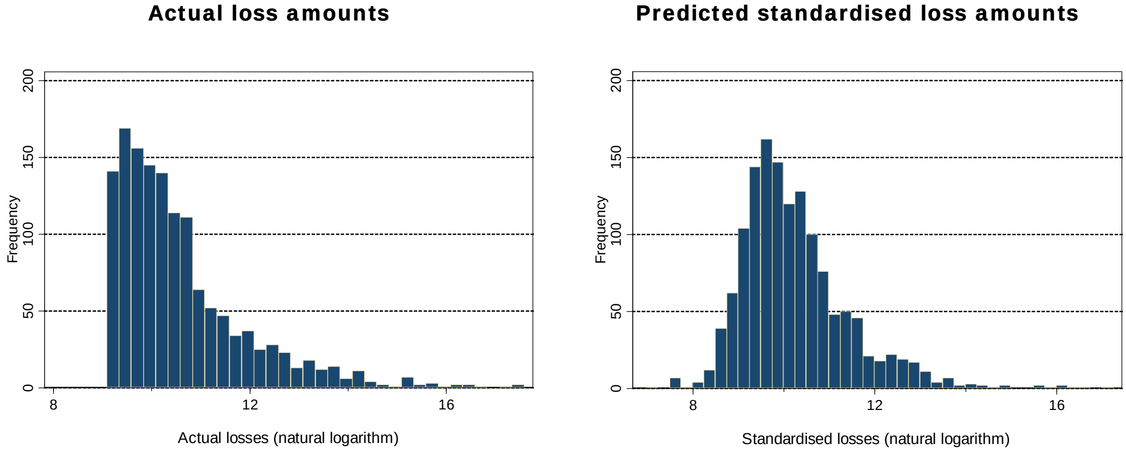

Standardised loss amounts are based on the model predictions and are explained by general factors of operational risks that are assumed to affect all firms symmetrically, irrespective of their actual sizes or control environments. This estimate represents a proxy for a "baseline" loss amount that is then scaled up or down to take into account idiosyncratic aspects of the insurer or event-specific features of the loss. As can be seen, the distribution of standardised loss amounts is less skewed than the distribution of actual losses although it is still affected by extremely large losses. The mean and median of the distribution of standardised loss amounts are £136,000 and £23,000, respectively.

The results of commonly used goodness-of-fit tests[7] suggested that the exponential distribution is most closely aligned with our sample of actual loss amounts. Since this analysis does not take into account any necessary scaling process, Figures 10 and 11 show empirical and fitted distribution functions for the severity of internal and scaled external losses occurring in two different insurers. The main difference between these firms is that Insurer A experiences substantially more loss events than Firm B.

The beta (generalised) distribution function offers the best fit to the internal loss data for firm A whereas the exponential function is more suitable to firm B's losses. If one looks at internal and scaled loss amounts, however, the model selection results are very different. We find the lognormal distribution function provides the best fit for both distributions, with the only difference that it has a heavier tail in the case of Insurer A.

[7] Specifically, we considered the Anderson-Darling, Kolmogorov-Smirnov and Chi-square tests. This family of statistical goodness-of-fit tests try to assess how close the theoretical model distribution is to the empirical distribution function.

Figure 10 Curve fitting for distribution of the size of losses – Insurer A

Internal losses only | **Internal and "scaled" external losses**

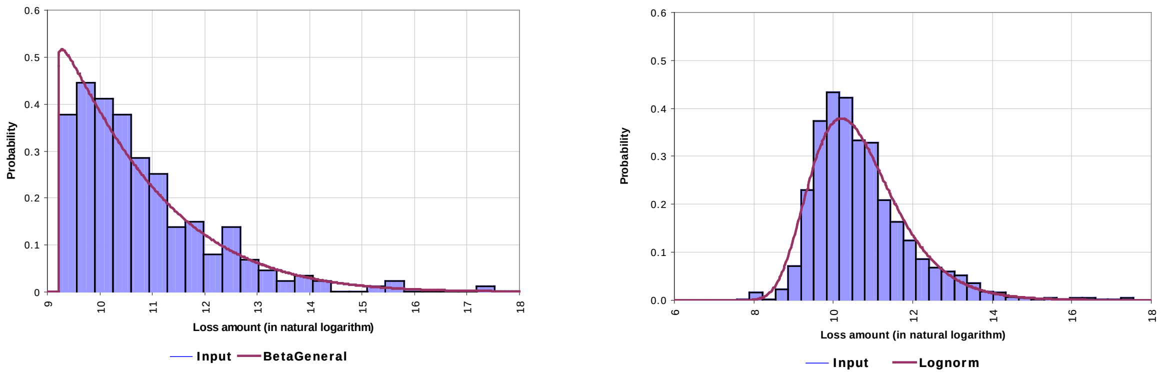

Note: This chart was created using Palisade's @Risk for Excel software.

Figure 11 Curve fitting for distribution of the size of losses – Insurer B

Internal losses only | **Internal and "scaled" external losses**

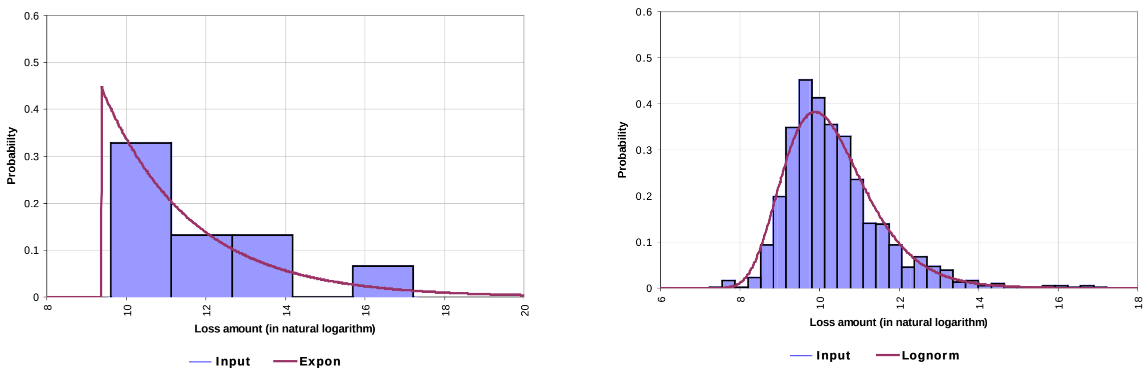

Note: This chart was created using Palisade's @Risk for Excel software.

Details of our scaling procedure can be found in Appendix A.1. In essence, we use the statistically significant parameters of our model to forecast values of the general and idiosyncratic factors for each loss event. The idiosyncratic predictions then allow us to translate all loss amounts in our sample into re-scaled losses that are more applicable to the profile of the insurer under consideration.

5.0 FREQUENCY ANALYSIS

This section deals with frequency distributions that describe the number of operational losses per unit of time. Particularly we look at the number of loss events reported per insurer per calendar quarter.

We focus on discrete probability distributions, or *probability functions*, which provide the probability for each possible outcome of a discrete random variable that may take values such as 0, 1, 2, 3, 4, etcetera. In our specific context, discrete distributions are defined for non-negative integers only, and represent the probability that certain number of losses or events causing operational losses occur in a given time period.

Section 5.1 sets out our analytical approach, Section 5.2 discusses a methodology for scaling the frequency of quarterly losses and lists our main findings, while Section 5.3 tackles the issue of curve fitting.

In summary, we find that:

- The size of the insurer, as measured by gross written premiums, is strongly and positively correlated with the number of losses. For a standard deviation rise in gross written premium, roughly £3.7bn, the projected number of operational losses per quarter increases by 24 per cent, holding other variables constant.

- The number of full-time employees does not appear to be a strong predictor for the frequency of losses. However, in the few specifications where we do observe a strong relationship, the correlation between full-time employees and number of operational losses is negative.

- Life business units are strongly correlated with higher frequencies of losses. Yet this could be partly driven by an overrepresentation of life firms in our sample.

- The estimated models account for 10% to 30% of the observed variability in the actual number of (quarterly) operational losses.

- Between the two models considered, the negative binomial regression provides a better fit to the observed frequency data than the Poisson regression model.

5.1 Analytical approach

We consider two frequency distributions that received a great deal of attention in the operational risk literature:

1. The Poisson probability distribution; and
2. The negative binomial distribution.

Their exact functional forms can be found the Appendix A.2.

The Poisson is one of the most popular in OR frequency estimation due to its simplicity and the fact it fits most loss databases very well (Cruz 2002). Its two basic underlying assumptions are that the probability of a loss event is the same for time intervals of equal length, and that the probability of a loss is independent across intervals.

As noted by Panjer (2006), the Poisson distribution has a number of useful modelling properties such as the fact that if operational losses follow a Poisson distribution and can be subdivided into different loss types or categories, then the distribution of the losses in the different categories is also Poisson but with a new parameter. This may be a useful feature when dealing with truncated data, because one can safely assume that the subset of losses will also be governed by a Poisson distribution function.

The negative binomial distribution has also been widely used in OR applications as an alternative to the Poisson distribution. It is also defined only for non-negative integers, but it offers more modelling flexibility than Poisson because it has more parameters. The negative binomial distribution is more suitable to datasets in which the variance of frequencies is larger than the mean, as the Poisson distribution assumes the mean and variance of the number of occurrences are equal (a property called "equidispersion").

5.2 Scaling model for the number of losses

In this section we consider a methodology for scaling frequency distributions in order to minimise biases that may arise when external and internal loss data are combined for economic capital estimations. As pointed out by Dahen and Dionne (2007), a relatively large literature has looked at ways to scale severity distributions but very few articles have so far attempted to develop similar approaches for the number of losses experienced by peer groups. Indeed, one could argue that scaling the frequency of external loss events is perhaps more critical than scaling their size, for the likelihood of an operational loss may be more sensitive to the insurer's scale of operations than its actual severity. In other words, the larger the insurer is, the bigger its exposures to operational risks are.

Our econometric analysis is based on models for count outcomes, where the explained variable is the number of losses experienced by 18 insurance companies per quarter in the period 2005-2008. These losses are in turn reported to ORIC on a quarterly basis. Since we expect the frequency of losses over a given time horizon will be linked to certain characteristics of the insurer such as size of operations or main business lines, we focus on the following explanatory variables (or exposure metrics):

- Average gross premium income (GWP);
- Average number of full-time employees (FTE); and
- Proportion of losses corresponding to the Life business unit (LIFE).

The first two variables capture the insurer's size while the last one represents a proxy for the business line. GWP and FTE have similar definitions to those provided above.

5.2.1 Findings of the scaling model for the frequency of losses

The principal findings of this part of our empirical analysis are set out below.[8]

Box 5.1 Results for frequency scaling

- While both Poisson and negative binomial models help us explain the observed variability in the number of operational losses, the latter offers the best fit.
- The size of the insurer is strongly correlated with the number of losses. We find that GWP is strongly and positively associated with the frequency of operational losses across all model specifications. However when FTE shows a statistically significant effect its sign is negative, meaning that the bigger FTE is, the smaller the number of looses is.
- As stressed in our severity analysis, the above results should be interpreted with care because GWP and FTE are positively correlated. Overall, evidence suggests that the insurer’s size is positively correlated with the number of loss events.
- To illustrate, our findings suggest that for a standard deviation increase in GWP, roughly £3.7bn, a firm’s mean number of quarterly operational losses increases by 24%, holding other variables constant.
- The percentage of losses associated with the life business unit has a significantly positive impact on the number of operational losses. For example, when the percentage of life losses rises from 0% to 47%, a firm’s mean frequency of quarterly losses increases by 127%, holding all other variables constant. We remain cautious about this finding, though, as it could be partly driven by some overrepresentation of life insurers in the ORIC database.
- Whereas the estimated models are highly significant, they only explain a small fraction of the variation in the observed number of operational losses – typically around 10% to 30% of the observed variability in quarterly losses.

Robustness tests

- We estimated an additional regression on a smaller sample size including only observations with strictly positive number of quarterly losses. This also allowed us to disregard cases where Life was zero simply because the firm reported no loss in that quarter. The sample shrank from 288 to 206 observations, but the sign and significance level of the coefficients did not change significantly.

[8] Detailed econometric results can be found in Appendix A1.

- A statistical test strongly suggested that the negative binomial regression model is preferred to the Poisson regression model because of the level of dispersion in the actual loss frequency distribution.

Figure 12 illustrates our key findings. The chart shows the predictions made by three fitted models: (1) the univariate Poisson distribution, (2) the Poisson distribution with the three explanatory variables, and (3) the negative binomial distribution with the same three regressors. Although the predictions of the multivariate Poisson regression model are an improvement over the univariate model, we find the negative binomial regression model provides the best statistical fit to the observed count data.[9]

Figure 12 Observed and predicted frequency of quarterly losses

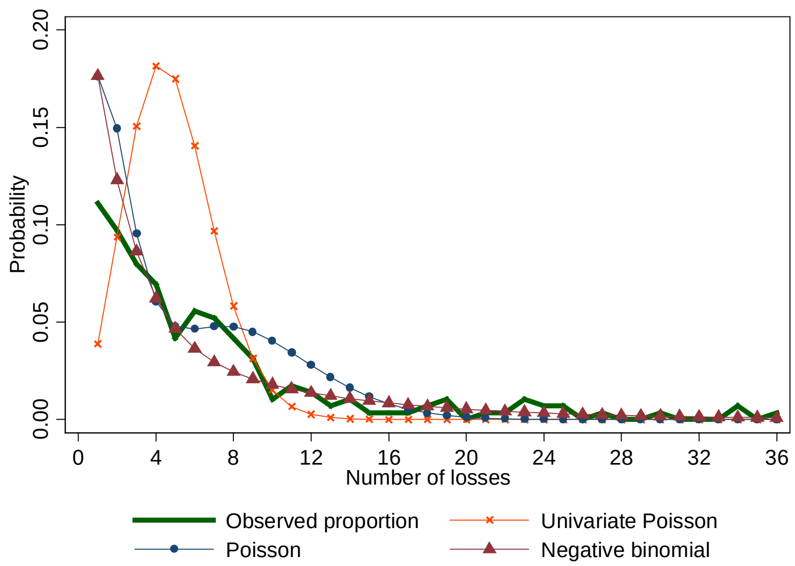

Note: For the univariate Poisson regression, the number of losses is regressed on a constant term only.
Source: ABI Research based on results of Poisson and negative binomial regression models.

5.3 Insurer's size and predicted number of losses

Figures 13 and 14 show, respectively, predicted probabilities for number of operational losses for the full sample and the subsample of observations with non-zero loss events in that quarter. The estimates are based on the negative binomial model. To illustrate the impact of the firm's size on the predicted frequency of losses, we focus on the 10% and 90% percentiles of the distribution of premiums, while holding all other variables

9 This likelihood-ratio test is automatically reported by Stata after the estimates of the parameters and differs from others in that its significance level is adjusted to account for underlying truncation in the sampling distribution. For more information, consult Stata's technical manuals.

constant at their average value, and consider the predicted probability and cumulative functions for number of losses. We concentrate on number of loss events in the range 0 to 9 since this captures the most important part of the distribution.

As can be observed, the insurer's size (as measured by GWP) has a material effect on the distribution of the predicted number of losses. The fitted curves differ not only in terms of expected values, but also in terms of aggregate shape. For example, insurers being at the lower end of the distribution of GWP are predicted to experience more than nine losses only 4 out of 100 times, whereas insurers in the upper end of the size distribution would experience more than nine losses 12 out of 100 times (Figure 13). If one considers instead the smaller sample of non-zero losses (Figure 14), our forecasts rise to 12% and 30% respectively.

Figure 13 Insurer's size and predicted probabilities for number of losses based on negative binomial model (full sample of losses)

PROBABILITY FUNCTION | **CUMULATIVE FUNCTION**

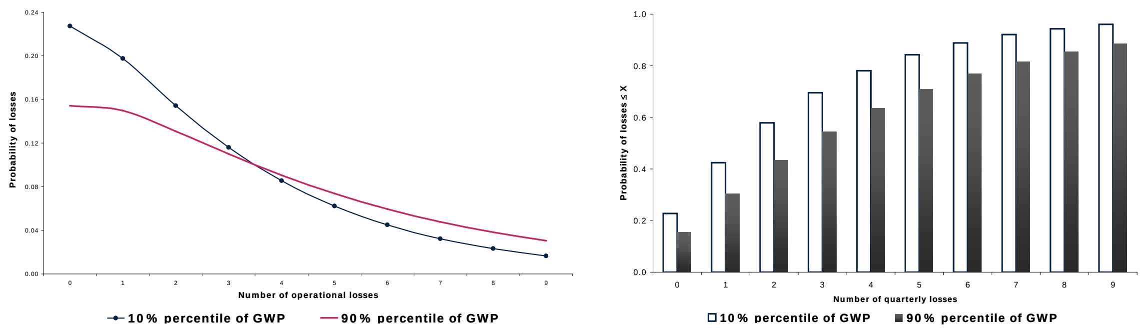

Source: ABI Research Department.

Figure 14 Insurer's size and predicted probabilities for number of losses based on negative binomial model (sample of non-zero losses)

PROBABILITY FUNCTION | **CUMULATIVE FUNCTION**

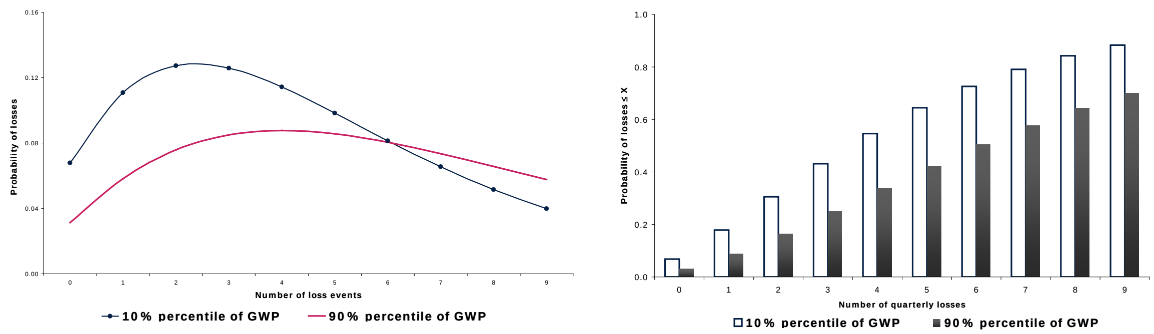

Source: ABI Research Department.

Another way to gauge the impact of individual scaling factors on the predicted number of operational losses is based on our Poisson regression model. In this case we use the estimated coefficients to compute the underlying parameter of the Poisson probability function for the 10% and 90% percentiles of the dispersion of premium income, while holding all the other variables at mean levels. These parameters were then plugged into a Poisson distribution, and the results are shown in Figure 15. The conclusions drawn are consistent with those obtained for the negative binomial model: the shape of the predicted probability distribution for the frequency of operational losses is very sensitive to the size of the insurer under consideration. These discrepancies can have sizeable implications for the calibration of capital models for operational risks.

Figure 15 Insurer's size and predicted probabilities for number of losses based on Poisson model (full sample)

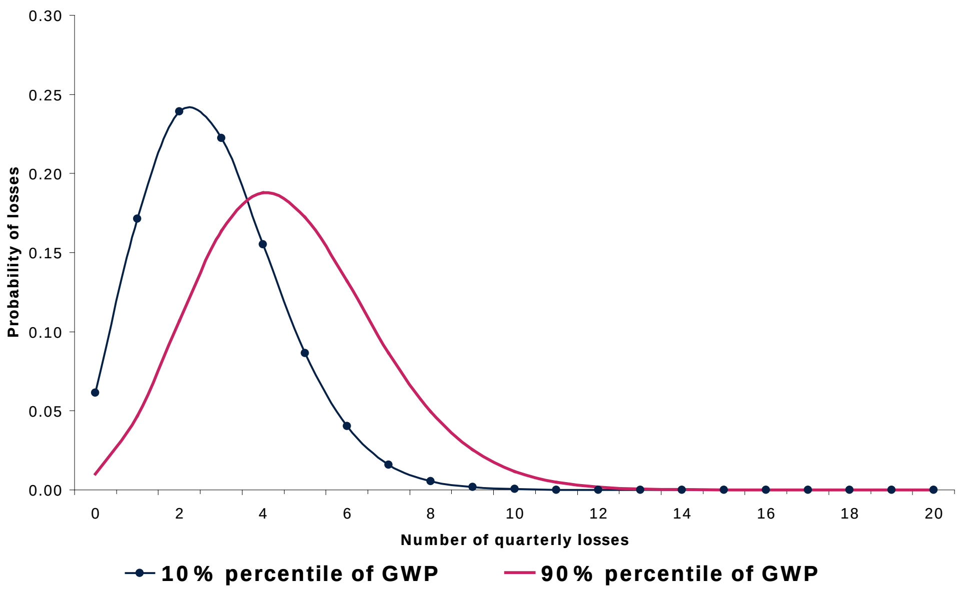

Source: ABI Research Department.

6.0 CONCLUDING REMARKS

This paper considered analytically sound methodologies to combine internal and scaled external loss data in order to build representative samples of historical operational loss events and ameliorate possible sample-selection biases. Among the main findings are that premium income is strongly and positively correlated with the size of operational losses, whereas the number of full-time employees is strongly and positively correlated with the number of losses in a given quarter. The size of recovery amounts after a loss event and some customer-related event types are also important to re-scale the severity of operational losses. The business line where losses occur in turn matters to explain the variability in observed frequency of loss events.

Scaling losses from external sources also affects markedly the results of curve-fitting exercises. For example, the exponential and beta generalised functions were the best distributions for the size of unscaled losses, but a lognormal distribution provided the best fit after scaled external losses are considered. Our scaling models also suggested the negative binomial model was a better fit to the observed number of losses than the Poisson model. Our results have important implications for the calculation of economic and regulatory capital models for operational risks in insurance.

Operational losses in insurance, particularly relatively severe ones, do not happen very often and also take time to crystallise. Therefore collecting and aggregating historical loss data are time-consuming and challenging tasks. While we checked the robustness of our findings, these results reveal emerging trends in the data that will need to be revisited as the ORIC database grows. This will provide a good opportunity to re-calibrate some of our main estimates. As the ORIC database continues to expand, we shall analyse these issues further and consider the question of how scaling models for the size and number of operational losses in insurance can be effectively combined in the context of internal models under Solvency II.

A1 ECONOMETRIC RESULTS

This appendix explains our econometric approach to scaling the severity and frequency of operational losses, and discusses the output of our estimations.

A1.1 Models for the size of losses

Our econometric approach for scaling the size of operational losses is based on linear regression techniques. Following previous studies focused on the banking sector [e.g. Na et al. (2006), Dahen and Dionne (2007), and Cope and Labbi (2008)], we assume the financial loss associated to operational failures can be partly explained by "general" and "idiosyncratic" factors. The size of the loss also depends on random factors, which for modelling purposes are encapsulated in a white-noise error term.

General elements capture features of the environment in which firms operate, such as prevailing macroeconomic conditions or the industry's regulatory framework, and are assumed to affect all insurers similarly. Other things affecting the size of losses that remain constant over the period will also be captured by this common component. Firm-specific elements are in turn allowed to vary across firms and loss events. They include things such as the size of the insurer, the lines of business affected, and loss event type categories.

Our econometric model assumes the following functional form for loss amounts (*L*):

$$L = A \times \exp\left(\sum \beta X + \delta Z\right),$$

where *A* captures the general aspect of the loss while *X* and *Z* refer to firm-specific and event-specific regressors, respectively. In this multiplicative model for loss amounts, a baseline amount is multiplied by a scaling factor that takes into account idiosyncratic features of the firm and characteristics of the loss event under consideration (e.g. size of the insurer, business lines). They are called "exposure metrics". The multiplicative and exponential assumptions made above have been widely used in the literature looking at ways of scaling operational losses in financial services.

After taking logarithms on both sides of the above expression, we get an equation for the size of losses that can be estimated using ordinary least squares (OLS) techniques:

$$Ln\, L_i = a + \sum b\, x_{ij} + \sum d\, z_{im} + e_i\,.$$

We fit this *semilog* model, where L stands for the gross loss per event in our ORIC sample, x and z are exposure variables (regressors), and e is an error term with the usual properties. The sign and significance levels of the coefficients b and d represent the direction and strength of any association between the exposure metric and the size of the loss, respectively. Note that $a = Ln\, A$.

The explanatory variables used in our linear regression models are the following:

- **Size of the insurer**. To control for the firm's size, we use average premium income (single and regular premiums, where appropriate) and average number of full-time employees (excluding contractors and outsourced activities) over the reporting period. These regressors, denoted by "GWP" and "FTE" respectively, are firm-specific and remain constant over the period of analysis.[10]

- **Firm-specific effects**. Our regressions include indicator variables (i.e. dummy variables) that identify the five top firms in terms of number of loss events in our ORIC sample. Each of these insurers has more than 100 loss amounts over the period. Our goal is to try to capture other firm-specific features that may be relevant to explain the size of losses, such as the insurer's control environment or quality of operational risk management.

- **Recovery**. The ORIC database also contains information on recovery amounts, that is, sums of money that the firm was able to recuperate after the loss event. Recovery amounts could for example be the result of insurance payouts, like for instance when physical assets are damaged or a building is flooded. We created an indicator variable (denoted "Recovery") that takes value 1 if the event had a positive recovery amount, and 0 otherwise.

- **Business line**. We expect main business lines to have a potential impact upon the severity of an operational loss. To capture this effect, we build an indicator variable (denoted "Life") that takes value 1 if the loss event occurred in the Life business unit and 0 otherwise.

- **Business function**. The losses originated in certain business functions may be higher than those of other functions, on average. To account for this potential driver of loss amounts, we use indicator variables that take value 1 if the loss occurred in certain business function and 0 otherwise. We constructed indicator variables identifying the following business functions: Customer Service/Policy Administration ("CS"), Sales and Distribution ("S&D"), and Claims ("Claims").

- **Loss event type**. Loss amounts in certain event types may be, on average, higher or smaller than losses in other categories. We use indicator variables that take value 1 if the loss falls into certain event type (at Level 2) and 0 otherwise. The particular event types we considered are Advisory Activities ("Advisory"), Customer Account Management ("Customer"), Systems ("Systems"), and Transaction Capture, Execution and Maintenance ("Transactions"), which are the most frequent (Level 2) categories in our ORIC sample.

[10] Econometric studies of operational losses in banks often use revenues or income to measure the size of institutions and/or business lines.

A region indicator wasn't feasible because most losses in our sample occurred in the UK. However, the geographical location of the loss may be an important predictor of its severity and so we believe this issue deserves more attention in future research.

Table 3 shows descriptive statistics for the dependent and explanatory variables of our OLS regressions. The majority (76%) of losses refer to the life business unit, whereas a high number (43%) of losses occurred in the Customer Service/Policy Administration function. Only 13% of the individual loss events had positive recovery amounts.

Table 3 Descriptive statistics

	Mean	Standard deviation	Min	Max
Loss amount (Ln of losses in £)	10.60	1.34	9.21	17.53
GWP (£bn)	4.64	4.68	0.33	14.64
FTE (thousands)	4.58	2.58	0.22	14.49
Recovery	0.13	0.33	0	1
Life	0.76	0.43	0	1
CS	0.43	0.50	0	1
S&D	0.21	0.41	0	1
Claims	0.08	0.28	0	1
Advisory	0.09	0.29	0	1
Customer	0.14	0.35	0	1
Systems	0.09	0.29	0	1
Transactions	0.39	0.49	0	1

Note: The underlying number of observations is 1,388.

Source: ORIC sample of losses.

A1.1.1 Econometric results

Table 4 presents the estimated coefficients of our OLS regression models to scale the size of operational losses.

Table 4 OLS model for the size of losses

	Model specification				
	(1)	(2)	(3)	(4)	(5)
GWP	0.15***	-0.002	0.02	0.03	0.03
FTE		0.18***	0.18***	0.16***	0.16***
Recovery			0.23*	0.32***	0.45***
Life			-0.22***	-0.10	-0.12
CS				-0.55***	-0.48***
S&D				-0.04	-0.16
Claims				-0.39***	-0.29**
Advisory					0.39**
Customer					-0.40***
Systems					0.14
Transactions					0.11
Constant	10.40***	10.08***	10.12***	10.30***	10.22***
R-squared	0.09	0.15	0.16	0.19	0.20

Note: The coefficients of the insurers' dummy variables are omitted from the output but were included in all model specifications. The variables GWP and FTE are expressed in £bn and '000s, respectively. All models include robust variance estimators and are statistically significant according to the usual F test. The symbols ***, ** and * indicate that the coefficient is statistically significant at the 1, 5 and 10 per cent levels, respectively.

Source: ORIC database.

Table 5 contains the results of additional OLS regressions that assess the robustness of our findings. Specifically we re-estimated model 5 above using net loss amounts as dependent variable (specification 6), we fitted a quantile regression model evaluated at the median (specification 7)[11], and we ran a linear regression model but accounting for the fact that loss amounts are truncated at the level of £10,000 (specification 8) – specifications 7 and 8 were estimated using maximum likelihood techniques.

[11] The objective of this quantile regression is to estimate the median of the dependent variable, rather than is mean, conditional on the values of the independent variables. This method puts less weight on extreme deviations from the regression line, since it minimises the sum of the *absolute* residuals rather than the sum of the *squares* of the residuals, as in OLS.

Table 5 Robustness tests on model for the size of losses

	Model specification		
	(6)	(7)	(8)
GWP	0.06**	0.001	0.28***
FTE	0.15***	0.23***	0.33***
Recovery	-0.94***	0.29**	1.72***
Life	-0.26**	-0.14	-0.87**
CS	-0.36***	-0.29***	-1.97***
S&D	-0.09	-0.04	-0.75
Claims	-0.27*	-0.36**	-1.15**
Advisory	0.37**	-0.11	1.64**
Customer	-0.49***	-0.44***	-2.23***
Systems	0.07	0.04	0.54
Transactions	0.03	-0.10	0.25
Constant	10.21***	9.85***	6.46***
Observations	1,271	1,388	1,350
R-squared	0.20	0.09	N/A

Note: The coefficients of the insurers' dummy variables have been omitted. The regressors GWP and FTE are expressed in £bn and '000s, respectively. All models include robust variance estimators and are statistically significant according to the usual F test. The symbols ***, ** and * indicate that the coefficient is statistically significant at the 1, 5 and 10 per cent levels, respectively.

Source: ORIC database.

The sign and significant level of the coefficients remain relatively stable across models. As expected, Recovery switches its sign when we look at net losses. Premium income seems important to explain the variability in net losses but not in gross loss amounts. The effect of the regressors on loss amounts also increases in the truncated regression (specification 8).

Table 6 shows results of quantile regressions estimated at the 25th, 50th (median) and 75th percentiles of the loss severity. GWP and FTE maintain their sign and significance levels. So too does Recovery. Customer-related business function and event type also remain relatively stable across quantiles. Only the coefficient on *Advisory* does lose its significance level when moving from the lower to the upper quantile of the loss severity distribution.

Table 6 Quantile regressions for the size of losses

	Percentile		
	25th	50th	75th
GWP	0.03	0.001	0.001
FTE	0.09***	0.23***	0.23***
Recovery	0.26***	0.29**	0.29**
Life	-0.02	-0.14	-0.14
CS	-0.17*	-0.29***	-0.29***
S&D	-0.01	-0.04	-0.04
Claims	-0.17**	-0.36**	-0.36**
Advisory	-0.13**	-0.11	-0.11
Customer	-0.26***	-0.44***	-0.44***
Systems	0.10	0.04	0.04
Transactions	-0.11	-0.10	-0.10
Constant	9.49***	9.85***	9.85***
Observations	1,388	1,388	1,388
R-squared	0.05	0.09	0.17

Note: The coefficients of the insurers' dummy variables have been omitted. The regressors GWP and FTE are expressed in £bn and '000s, respectively. The variance-covariance matrix of the estimators was obtained by bootstrapping and all models are statistically significant. The symbols ***, ** and * indicate that the coefficient is statistically significant at the 1, 5 and 10 per cent levels, respectively.

Source: ORIC database.

A1.1.2 Scaling external losses

A key feature of the above econometric model is that it allows us to "scale" external loss amounts. Effectively it provides us with a normalisation procedure to translate loss events occurred in other firms into equivalent losses that could have been experienced by the firm under consideration. Both internal and external operational loss events are thus expressed in a common metric.

Since the equation

$$Ln\,L_i = a + \sum b\,x_{ij} + \sum d\,z_{im}$$

is assumed to apply to all losses/insurers in the database, we can use the fact that the general component a is constant across insurers to find the value of scaled losses for insurer A (i.e., B's losses translated into A's losses) as follows:

$$Loss_A = \frac{\hat{h}_A}{\hat{h}_B} \times Loss_B\,.$$

Hence the scaling factor is equal to the ratio of the estimated idiosyncratic components for the two firms being compared. The function $\hat{h}_i$ is built using all the statistically significant regressors of our baseline OLS model:

$$\hat{h}_i = b_1\, GWP_i + b_2\, FTE_i + d_1 reco_i + d_2\, CS_i + d_3\, SD_i + d_4\, claims_i + d_5\, adv_i + d_6\, custom_i$$

We also included the insurer's dummy variable if it was applicable and significant. We used this method to calculate the scaled losses discussed in the main text.

A1.2 Models for the number of losses

Applying linear regression techniques to count outcomes may result in inefficient, inconsistent, and biased estimates. To avoid this problem, in this paper we focus on two alternative regression models that are specifically designed to account for discrete outcomes of the dependent variable:

- Poisson regression, and
- Negative binomial regression.

Our count variable of interest is the number of individual operational losses reported to ORIC by a firm in a given calendar quarter, whereas the explanatory variables are the firm's average gross written premiums (denoted by "GWP") and average full-time employees (denoted by "FTE") over the period, and the proportion of quarterly losses corresponding to the Life business unit (denoted by "Life"). The definitions of GWP and FTE are similar to the ones provided above. The time horizon refers to the quarter in which the loss actually occurred, as opposed to the quarter in which it was reported to the consortium. In those quarters where an insurer did not report losses, the observed frequency is zero.

Table 7 contains basic descriptive statistics for the variables included in our models.

Table 7 Descriptive statistics

	Mean	Standard deviation	Min	Max
Number of events per quarter	4.82	6.51	0	36
GWP (£bn)	3.26	3.70	0.33	14.64
FTE (thousands)	4.25	3.99	0.22	14.49
Life	0.42	0.47	0	1

Note: The underlying number of observations is 288.
Source: ORIC database.

We have 288 observations because we track 18 firms over a time period of 16 quarters (2005:Q1 to 2008:Q4). Note that as each company contributes exactly one frequency data point every quarter, the descriptive statistics for GWP and FTE reported above are different from the ones presented in the analysis of severity of operational losses.

The econometric outputs reported throughout this section were estimated using the econometric package Stata 9.2. Maximum likelihood methods were used to compute the relevant parameters. A review of applications of regression models for count data can be found, for example, in Long and Freese (2003).

A1.2.1 Poisson regression

The Poisson regression model specifies that each observed count y_i is drawn from a Poisson distribution with parameter μ>0, which is in turn related to the explanatory variables x_i. The main equation of the model is

$$\Pr(y \mid \mu) = \frac{\mu^y e^{-\mu}}{y!} \qquad for\ y = 0,1,2,\ldots$$

The Poisson regression model allows each observation to have potentially a different value of μ. In other words, it assumes that the observed count is drawn from a Poisson distribution with mean μ_i, where μ_i is estimated from observed characteristics. The most common formulation in the literature is to capture the observed heterogeneity in the data through the following functional form:

$$\mu_i = E(y_i \mid x_i) = \exp\left(\sum \beta_i x_i\right).$$

Note the exponential function forces μ to be positive, which in this context is necessary because count outcomes cannot be negative. More technical details of the Poisson regression model can be found in Greene (2000) and Cameron and Trivedi (1998).

Econometric results

Table 8 shows the estimated coefficients of our Poisson regression model.

Table 8 Poisson regression model (PRM)

	Model specification					
	(1)	(2)	(3)	(4)	(5)	(6)
GWP		0.08***	0.09***	0.07***		0.07***
FTE			-0.03*	-0.03*	0.01	-0.04**
GWP/FTE					0.14	
Life				1.65***	1.71***	0.93***
Constant	1.57***	1.26***	1.34***	0.44***	0.37***	1.16***
Observations	288	288	288	288	288	206
R-squared	<0.01	0.06	0.06	0.31	0.28	0.17

Note: The regressors GWP and FTE are expressed in £bn and '000s, respectively. All models include robust variance estimators and are statistically significant according to the usual Wald test. The symbols ***, ** and * indicate that the coefficient is statistically significant at the 1, 5 and 10 per cent levels, respectively.

Source: ORIC database.

In model 5 we included a variable constructed as GWP over FTE and dropped premium income from the regression because it was highly correlated with the new regressor. It deserves noting that specification 6 has a smaller sample size because the underlying observations are only those where the number of quarterly losses is strictly positive. This was done to exclude cases where Life takes value zero simply because the insurer had no loss in that quarter. The coefficients of the regressors do not change much.

A1.2.2 Negative binomial regression

One serious shortcoming of the Poisson regression model is the assumed equality of the conditional mean and variance functions. Thus the model often underestimates the actual amount of dispersion in the observed outcomes. An alternative specification that does allow for heterogeneity among observations and overdispersion is based on the negative binomial model, which is obtained by adding an individual, unobserved effect to the conditional mean.[12] Specifically,

$$\begin{aligned}\tilde{\mu}_i &= \exp\left(\sum \beta_i\, x_i\right)\exp(\varepsilon_i) \\ &= \mu_i\, \delta_i\end{aligned}$$

where $\delta \equiv \exp(\varepsilon)$. A gamma distribution for the individual effect ε and the condition $E(\delta) = 1$ are assumed to facilitate model identification. It can then be shown that this normalisation leads to a density function for the count outcome that is one form of the **negative binomial distribution** with mean μ – see Long (1997) and Greene (2003).

12 Dionne and Vanasse (1992) discuss evidence on the superiority of the negative binomial regression model over the Poisson regression model in the context of automobile accidents.

Econometric results

Table 9 shows the estimated coefficients of our negative binomial model.

Table 9 Negative binomial regression model

	(1)	(2)	(3)	(4)	(5)	(6)
GWP		0.08***	0.09***	0.06***		0.06***
FTE			-0.02	-0.03	0.01	-0.04**
GWP/FTE					-0.01	
Life				1.77***	1.84***	0.98***
Constant	1.57***	1.27***	1.34***	0.40***	0.42**	1.16***
Observations	288	288	288	288	288	206
R-squared	<0.01	0.01	0.01	0.09	0.09	0.06

Note: The regressors GWP and FTE are expressed in £bn and '000s, respectively. All models include robust variance estimators and are statistically significant according to the usual Wald test. The symbols ***, ** and * indicate that the coefficient is statistically significant at the 1, 5 and 10 per cent levels, respectively.

Source: ORIC database.

A1.2.3 Comparison of econometric results

Table 10 compares the estimated coefficients and respective z-values of specification 4 for the Poisson regression (PR) and negative binomial regression (NBR) models. As can be seen, regression results are highly aligned in terms of both sign and significance levels of the coefficients, with the only exception that FTE is marginally insignificant in the NBR model.

Table 10 Coefficients of the PR and NBR models

	PRN		**NBRM**	
	Coefficient	z-value	Coefficient	z-value
GWP	0.07	3.18	0.06	3.37
FTE	-0.03	-1.83	-0.03	-1.43
Life	1.65	11.56	1.76	11.49
Constant	0.44	3.42	0.40	2.75

In the presence of overdispersion, estimates from the PR model are inefficient with standard errors that are biased downward. Therefore we also carried out a likelihood-ratio test for overdispersion in the data, which suggested that the negative binomial regression model is preferred to the Poisson regression model.

A2 PROBABILITY DISTRIBUTIONS FOR SEVERITY AND NUMBER OF LOSSES

Once external loss events have been scaled, it is important to assess which statistical distribution provides the best fit to the internal and scaled external losses. In effect the fitted severity and frequency loss distributions are key inputs into the predictions of stochastic simulations that ultimately feed the firm's economic and regulatory capital models. In this appendix we explain briefly the properties and functional forms of the statistical distributions for the size and number of losses considered in the paper.

A2.1 Distributions for the size of losses

A2.1.1 Beta generalised probability function

This distribution can be directly derived from the Beta distribution,

$$f(x) = \frac{x^{\alpha_1 - 1}(1-x)^{\alpha_2 - 1}}{B(\alpha_1, \alpha_2)},$$

by scaling the [0, 1] range of the Beta distribution with the use of a minimum and maximum value to redefine the range. B is the Beta function and the alphas are parameters of the distribution. The generalised Beta distribution has the following density function:

$$f(x) = \frac{(x - \min)^{\alpha_1 - 1}(\max - x)^{\alpha_2 - 1}}{B(\alpha_1, \alpha_2)(\max - \min)^{\alpha_1 + \alpha_2 - 1}}.$$

The shape parameters α_i are both positive, and the distribution is continuously defined in the range $\min \le x \le \max$.

A2.1.2 Exponential probability function

This distribution is the continuous time equivalent to the geometric distribution in the discrete time case. It represents the waiting time for the first occurrence of a process that is continuous in time and of constant intensity, and it has been applied in a wide variety of statistical procedures (e.g. queuing, breakdown modelling). It is also very popular in the field of operational risk.

The exponential distribution has the density function:

$$f(x) = \frac{e^{-x/\lambda}}{\lambda},$$

where $\lambda > 0$ is a scale parameter that equals the conditional mean of the distribution. As this conditional mean is constant, the expected size of the excess operational loss above certain threshold doesn't depend on the value of the threshold.

A2.1.3 Lognormal probability function

This distribution is a transformation of the normal distribution as it assumes that the natural logarithm of the underlying random variable is normally distributed. The log-normal density function thus describes the distribution of a variable whose logarithm is normally distributed, and is defined as follows:

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(\ln x - \mu)^2 / 2\sigma^2}$$

where x is the random variable denoting the size of the loss, $f(x)$ is the probability density function, μ denotes the mean, σ is the standard deviation, $\pi = 3.14159$ and $e = 2.71828$.

A2.1.4 Pareto probability function

The Pareto distribution is characterised by the following density function

$$f(x) = \frac{\theta \alpha^{\theta}}{x^{\theta+1}}.$$

This function is only defined for values greater than or equal to α. The parameters α and θ are positive and denote the scale and shape shifters, respectively. Having two parameters offers in principle more degree of modelling flexibility, albeit in applications the value of α is known and therefore set in advance. This leaves only one parameter to choose from.

A2.2 Distributions for the number of losses

A2.2.1 Poisson probability distribution

The Poisson probability function is defined as follows:

$$f(x) = \frac{\mu^x e^{-\mu}}{x!} \qquad for\ x = 0,1,2,\ldots$$

where x is a random variable denoting the number of times an event did occur, $f(x)$ is the probability of observing x operational losses, $\mu > 0$ is the expected number of losses over the specified time interval, and $e = 2.71828$. For the Poisson probability function, the mean and the variance of the number of occurrences are equal (this is often called "*equidispersion*").

The Poisson probability distribution has two main properties:

1. The probability of an occurrence is the same for any two intervals of equal length;

2. The occurrence or non-occurrence in any specified interval is independent of the occurrence or non-occurrence in any other interval.

A2.2.2 Negative binomial probability distribution

The probability function of the negative binomial distribution is given by:

$$f(x) = \binom{s+x-1}{x} p^s (1-p)^x,$$

where $s \geq 0$ represents the number of successes, *p* is the probability of success on any one trial, and the binomial coefficient is defined as

$$\binom{n}{x} = \frac{n!}{x!\,(n-x)!}.$$

The geometric probability distribution is the special case of the negative binomial distribution when $s = 1$. As can be easily seen, for this case the probability function boils down to

$$f(x) = p\,(1-p)^x.$$

When $s < 1$ the negative binomial distribution has a heavier tail (decays more slowly) than the geometric distribution, whereas if $s > 1$ the negative binomial distribution has a lighter tail than the geometric distribution.

REFERENCES

Cameron, A. C. and Trivedi, P. (1998): *The Analysis of Count Data*, Cambridge University Press, USA.

CEIOPS (2008): “Quantitative Impact Study 4 – Summary results and main messages”, *Stakeholder Meeting, October 2008, Germany*.

Cope, E. and Labbi, A. (2008): “Operational loss scaling by exposure indicators: evidence from the ORX database”, *Journal of Operational Risk, Vol. 3(4), pp. 25-45*.

Cruz, M. (2002): *Modeling, Measuring and Hedging Operational Risk*, John Wiley & Sons Ltd, UK.

Dahen, H. and Dionne, G. (2007): “Scaling models for the severity and frequency of external operational loss data”, *Working Paper 07-01, Canada Research Chair in Risk Management*.

De Fontnouvelle, P., DeJesus-Rueff, V., Jordan, J. and Rosengren, E. (2003): “Using Loss Data to Quantify Operational Risk”, *Federal Reserve Bank of Boston, Mimeo*.

Dionne, G. and Vanasse, C. (1992): “Automobile Insurance Ratemaking in the Presence of Asymmetrical Information”, *Journal of Applied Econometrics, Vol. 7(2), pp. 149-165*.

Doff, R. (2007): *Risk Management for Insurers. Risk control, Economic Capital and Solvency II*, Risk Books, Incisive Financial Publishing Ltd, UK.

FSA (2008a): “Insurance Risk Management: The Path to Solvency II”, *FSA Discussion Paper 08/4, September 2008, UK*.

FSA (2008b): “FSA UK country report. The fourth Quantitative Impact Study (QIS4) for Solvency II”, *December 2008, UK*.

Greene, W. H. (2000): *Econometric Analysis*, 4th Edition, Prentice-Hall, USA.

Long, J. (1997): *Regression Models for Categorical and Limited Dependent Variables*, Advanced Quantitative Techniques in the Social Sciences Series 7, SAGE Publications Ltd, UK

Long, J. and Freese, J. (2003): *Regression Models for Categorical Dependent Variables Using Stata*, Stata Press, USA.

Na, H.S, van den Berg, J., Miranda, L. and Leipoldt, M. (2006): “An econometric model to scale operational losses”, *Journal of Operational Risk, Vol. 1(2), pp. 11-31*.

Panjer, H. (2006): *Operational Risk. Modeling Analytics*, John Wiley & Sons Inc., USA.

Sandström, A. (2009): *Handbook on Solvency for Actuaries and Risk Managers: Theory and Practice*, Chapman & Hall, Forthcoming.

Aims and scope: The Association of British Insurers (ABI) is the trade body representing the UK's insurance industry. The ABI Research Paper series is used to publish the research that the ABI carries out on behalf of its members in order to help inform the insurance industry and contribute to public policy debate.

Series Editor: Rebecca Driver, Director of Research and Chief Economist, ABI

Author: This paper was written by Mariano Selvaggi of the ABI Research Team and ORIC.

ABI Contacts: Copies of ABI Research Papers are available on the ABI website at: http://www.abi.org.uk/Bookshop/default.asp

Copies of ABI Research Papers may also be obtained from Research Department, Association of British Insurers, 51 Gresham Street, London, EC2V 7HQ; Tel: +44 (0)20 7216 7390; Fax: +44 (0)20 7216 7449; email: research@abi.org.uk.

For Press queries, please contact the ABI's media team on Tel: +44 (0)20 7216 7394; email: info@abi.org.uk.

ISBN

For more information, contact:

Association of British Insurers
51 Gresham Street
London EC2V 7HQ
020 7600 3333

www.abi.org.uk