




Investing in Insurance Risk

Insurance-Linked Securities
— A Practitioner's Perspective

BY ALEX KRUTOV



Investing in Insurance Risk

Investing in Insurance Risk

Insurance-Linked Securities – A Practitioner's Perspective

By Alex Krutov



Published by Risk Books, a Division of Incisive Financial Publishing Ltd

Haymarket House
28–29 Haymarket
London SW1Y 4RX
Tel: +44 (0)20 7484 9700
Fax: +44 (0)20 7484 9797
E-mail: books@incisivemedia.com
Sites: www.riskbooks.com
www.incisivemedia.com

© 2010 Incisive Media

ISBN 978-1-904339-56-4

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

Publisher: Nick Carver
Managing Editor: Jennifer Gibb
Development Editor: Sarah Hastings
Copy Editor: MFE Editorial Services
Typeset by Mark Heslington Ltd, Scarborough, North Yorkshire

Printed and bound in the UK by Printondemand-worldwide

Conditions of sale

All rights reserved. No part of this publication may be reproduced in any material form whether by photocopying or storing in any medium by electronic means whether or not transiently or incidentally to some other use for this publication without the prior written consent of the copyright owner except in accordance with the provisions of the Copyright, Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licensing Agency Limited of 90, Tottenham Court Road, London W1P 0LP.

Warning: the doing of any unauthorised act in relation to this work may result in both civil and criminal liability.

Every effort has been made to ensure the accuracy of the text at the time of publication, this includes efforts to contact each author to ensure the accuracy of their details at publication is correct. However, no responsibility for loss occasioned to any person acting or refraining from acting as a result of the material contained in this publication will be accepted by the copyright owner, the editor, the authors or Incisive Media.

Many of the product names contained in this publication are registered trade marks, and Risk Books has made every effort to print them with the capitalisation and punctuation used by the trademark owner. For reasons of textual clarity, it is not our house style to use symbols such as TM, ®, etc. However, the absence of such symbols should not be taken to indicate absence of trademark protection; anyone wishing to use product names in the public domain should first clear such use with the product owner.

While best efforts have been intended for the preparation of this book, neither the publisher, the author nor any of the potentially implicitly affiliated organisations accept responsibility for any errors, mistakes and or omissions it may provide or for any losses howsoever arising from or in reliance upon its information, meanings and interpretations by any parties.

Contents

About the Author	xiii
Preface	xv
 PART I: INTRODUCTION TO INVESTING IN INSURANCE RISK	 1
1 Investing in Insurance Risk	3
Investing in risk	3
Insurance risk	4
Insurance markets	4
Securities issued by insurance companies	6
Insurance-linked securities	7
Investing in insurance risk	9
2 Insurance-Linked Securities	13
Insurance-linked securities defined	13
Types of insurance-linked securities	14
Yield and diversification offered by insurance-linked securities	17
Market dynamics	19
 PART II: INVESTING IN AND MODELLING SECURITIES LINKED TO PROPERTY AND CASUALTY RISK	 21
3 Property Catastrophe Bonds	23
Securitisation of property insurance risk	23
Motivation for transferring natural catastrophe risk to the capital markets	24
Historical perspective	25
Risk transfer in insurance	26
Catastrophe bond structure	28
Default triggers	31
Number and types of perils	34

Term	35
Quantitative analysis	35
Investment performance of cat bonds	42
Market stability and growth	46
More on the sponsor and investor perspectives	47
Modelling property catastrophe insurance risk	51
Trends and expectations	52
4 Modelling Catastrophe Risk	55
The challenge of modelling catastrophe events	55
Importance of catastrophe modelling to investors	56
Modelling catastrophe insurance risk of insurance-linked securities	57
The science of catastrophes	57
Earthquake frequency and severity	59
Earthquake location	61
More on earthquake modelling	65
Tsunamis	69
Hurricanes	70
Historical frequency of hurricanes threatening the US	74
Seasonality of the hurricane risk in insurance-linked securities	77
Landfall frequency in peak regions	79
Hurricane frequency effects over various time horizons	81
Investor views on macro-scale frequency effects	83
Evolution of investor views on catastrophe modelling	86
Elements of hurricane modelling	88
Damage modelling	94
Financial loss modelling	95
Catastrophe model structure	97
Modelling terrorism risk	98
Modelling pandemic flu risk	101
Practical modelling of catastrophe risk	103
Data quality	106
Investor and catastrophe modelling	108
Catastrophe bond remodelling	109
Hurricane forecasting	110
Climate change	111
Sponsor perspective on modelling	112

Modelling as a source of competitive advantage to investors	113
Modelling as a source of competitive disadvantage to investors	115
Trends and expectations	116
5 Catastrophe Derivatives and ILWs	119
Index-linked contracts	119
Role of an index	120
Catastrophe derivatives defined	120
Industry loss warranties defined	122
Market size	122
Key indexes	
	Modelling industry losses
	The ILW market
	ISDA US wind swap confirmation text
	IFEX catastrophe derivatives
	CME hurricane derivatives
	Eurex hurricane futures
	More unusual products
	Comments on pricing
	Credit risk
	Basis risk
	The use of transformers
	Investor universe
	Mortality and longevity derivatives
	Investor and hedger perspectives
	Trends and expectations
6 Reinsurance Sidcars and Securitisation	
	Securitisation of reinsurance
	Reinsurance sidcars
	Sidcar structure
	Investor perspective
	Sponsor perspective
	Sidcar types
	Investor universe
	Considerations in investment analysis
	Trends and expectations

7	Credit Risk in Catastrophe Bonds and Other ILS	167
	Credit risk	167
	Credit risk and ILS	169
	Traditional solutions	169
	The need for new solutions	170
	Solutions to credit risk issues in insurance-linked securities	171
	Triparty repo arrangement	173
	Customised puttable notes	176
	Use of US Treasury money market funds as collateral	176
	Collateral options in collateralised reinsurance	177
	Trends and expectations	178
8	Weather Derivatives	181
	The broader definition of insurance-linked securities	181
	Weather derivatives defined	181
	Heating and cooling degree days	183
	Other types of weather derivatives	184
	Payout on standard options	186
	Exchange-traded weather derivatives	187
	Pricing models for weather derivatives	188
	Practical challenges in pricing	189
	Investing in weather derivatives	191
	Emissions trading	193
	Trends and expectations	194
	PART III: SECURITIES LINKED TO VALUE-IN-FORCE MONETISATION AND FUNDING REGULATORY RESERVES	197
9	Funding Excess Insurance Reserves	199
	Excess insurance reserves	199
	Some examples	199
	“Excess” reserves	201
	Funding solutions	201
	Embedded value and value-in-force securitisation	202
	Market fluidity	203
	RBC requirements leading to “unnecessary” capital strain	203
	Regulation XXX reserve funding	204
	Letter-of-credit facility for funding regulation XXX reserves	206

Securitisation of Regulation XXX reserves
Other solutions

207

Additional considerations for investors
Funding AXXX reserves
Loss portfolio transfer
Conclusion

10 Embedded Value Securitisation

Rationale for embedded value securitisation
Embedded value and value-in-force definitions
Direct monetisation versus true securitisation
Closed block
Investor perspective
Specific structures
Modelling
Stress scenarios
Ratings of EV securitisations
Examples of EV securitisation
Gracechurch/Barclays EV securitisation
Trends and expectations

PART IV: INVESTING IN AND MODELING LINKED TO MORTALITY AND LONGEVITY

11 Securitisation of Extreme Mortality Risk

The risk of extreme mortality
Securitisation of extreme mortality risk
The groundbreaking Vita securitisation
Other securitisations of extreme mortality
Basis risk
Credit enhancement
Investor types
Extreme mortality risk quantification and modelling
Current modelling approaches
Mortality derivatives
Additional considerations for investors
Trends and expectations

12 Life Insurance Settlements

Insurance policy as a tradable asset

Life settlements	264
Life settlement securitisations	266
Legal and ethical issues	267
Market participants	269
Current and future market size	269
Regulatory issues	272
The link between investor risk and consumer protection	273
Tax issues	274
Insurable interest	275
Investor- or stranger-originated life insurance policies	278
Contestability	280
Trust structures and investor due diligence	281
The use of not-for-profit organisations in life settlements	282
Investor perspective	285
Insurance industry perspective	289
Risks to insurers	290
Conclusion	293
13 Mortality and Longevity Models in Insurance-Linked Securities	295
Mortality and longevity	295
Mortality rates	296
Mortality tables	299
Population mortality tables	301
Mortality dynamics	306
Select and ultimate tables	308
Credibility theory approach	310
Longevity improvements	312
Lee–Carter and related methods	315
Markov process of mortality and morbidity	316
Direct age transform in mortality modelling	319
Mortality and longevity shocks	320
Conclusion	321
14 Valuation of Life Settlements and Other Mortality-Linked Securities	323
Modelling investment performance of life settlements	323
Life expectancy	325
Methodology changes in the calculation of life expectancy	327

Underwriting concepts
Debits

328

Choice of mortality table
2008 valuation basic table
Relative risk ratios
Underwriting for older ages
Choosing the LE
LE shopping
Assumed premiums
Being paid for the risk
Conclusion

15 Longevity Risk Transfer and Securities

Longevity risk
Need to transfer longevity risk
Longevity improvements
Natural hedges
Primary mechanisms of longevity
Longevity swaps
Mortality forwards and survival
Longevity bonds
More on other solutions for longevity
 in a DB pension fund
Indexes of longevity
Investors in longevity
Market developments
Extension risk in traded policies
Trends and expectations

PART V: MANAGING PORTFOLIO RISK

16 Managing Portfolios of Catastrophe Risk

Portfolio construction
Exotic beta
How catastrophe risk is different
Measures of return and risk
Managing a portfolio of catastrophe risk
 company

Managing a portfolio of catastrophe insurance-linked securities	407
Instrument types	409
Portfolio constraints	411
Standard tools and the modelling of individual securities	413
Portfolio optimisation	416
Pitfalls of standard optimisation techniques	426
Remodelling and portfolio optimisation	427
Sensitivity analysis and scenario testing	428
Additional considerations	429
Performance measurement	431
Conclusion	432
17 Managing Portfolios of Multiple Types of ILS	435
Types of insurance-linked securities	435
Rationale for combining different types of ILS in the same portfolio	437
Correlation among different types of ILS	438
Tenor and liquidity	438
Portfolio optimisation	439
The argument against combining ILS of multiple types in the same portfolio	441
Portfolio valuation issues	441
Performance measurement	442
Investment management policy	446
Risk management	446
Conclusion	448
18 Conclusion	449
References	453
Index	469

About the Author

Alex Krutov is managing director of Century Atlantic Capital Management, where he developed an investment strategy across all types of insurance-linked securities (ILS) and collateralised reinsurance, as well as portfolio optimisation and risk management techniques for ILS and reinsurance. Prior to joining the firm, he was president of Navigation Advisors LLC, a New York management-consulting firm focused on the insurance industry, capital markets, and general management. Prior to founding Navigation Advisors, Alex was employed in a variety of roles, including officer-level positions, at companies such as Transatlantic Reinsurance Company, American International Group (AIG), Reliance Group, UBS Warburg, and AXA Financial.

Alex's primary expertise and experience involve the products that bridge the gap between (re)insurance and capital markets. He has strong expertise in insurance securitisation, alternative risk transfer, reinsurance and insurance underwriting, portfolio issues in investing in insurance-linked securities, risk analysis, pricing of catastrophe (re)insurance risk, and general management.

Alex is a member of the American Academy of Actuaries, the Casualty Actuarial Society, and the Society of Actuaries. He chairs the Risk-Based Capital Committee of the American Academy of Actuaries. In addition to his actuarial credentials, Alex Krutov holds an MBA in management and finance from the Columbia University Graduate School of Business. He also holds an MS in physics.

Alex Krutov can be reached at alex.krutov@gmail.com or through the publisher.

Preface

The title of this book, *Investing in Insurance Risk*, might sound strange to an investor unfamiliar with securities linked to insurance. Any investment involves risk, so investors are not averse to accepting it; but risk is not what we generally want to invest in. We want to invest in securities that will likely generate healthy returns. When investing in a security, we pay for its probabilistically distributed future return. The uncertainty associated with the investment return, including the chance of the return being negative, is the risk we assume. In fact, many investors actively seek risky assets to invest in, as long as they believe they will be properly compensated for assuming those risks.

While risk is an integral part of investing, we generally do not think in terms of “investing in risk.” For securities whose performance is directly linked to insurance risk, however, the focus on the risk is so great and its nature so unusual, that it does make sense to speak in terms of investing in insurance risk. The insurance industry is concerned with measuring and managing risk, and so are the investors in securities with embedded insurance risk.

INSURANCE-LINKED SECURITIES

Any investment in traditional securities – such as common stock or bonds – of an insurance or reinsurance company may be seen as an investment in insurance risk, if we define insurance risk as simply any risk to which insurance companies are exposed. In addition, insurance securitisation has created a new asset class – referred to as insurance-linked securities (ILS) – that affords investors exposure to a more “pure” form of insurance risk. Examples include the risks of catastrophic insured losses, from hurricanes and earthquakes to those resulting from spikes in mortality rates due to pandemic events. This type of risk does not have to be associated with a catastrophic event, though; potential improvements in human longevity, for example, could have a severe financial impact on insurance companies selling annuity products. Longevity improvements are not a catastrophic event *per se*, but the financial consequences can be catastrophic. Such risks, although labelled

insurance risks, do not have to originate in the insurance industry. Pension plans are even more exposed to longevity risk than insurance companies. Furthermore, insurance risk transferred to the capital markets does not have to involve any catastrophic component at all; as is the case when an insurance company transfers some of its pure insurance risk to investors simply to use its capital more efficiently or to reduce earnings volatility.

Insurance risk lacks a clear, unambiguous definition, and so do insurance-linked securities. The best known types of insurance-linked securities – catastrophe bonds and life insurance settlements – are clearly in the ILS category, but some others, such as weather derivatives or collateralised reinsurance, can reasonably be seen as not belonging to this asset class.

Insurance-linked securities include a number of risks that can be highly correlated to traditional financial assets. At the same time, however, the “pure” insurance risk typically has low correlation with the rest of the capital markets. This low correlation is one of the primary reasons investors have been watching this asset class with interest. The overall degree of correlation of insurance-linked securities with the markets can vary; for example, the correlation of properly structured catastrophe bonds is much lower than that of embedded value securities.

INSURANCE INDUSTRY

Even though not all “insurance risk” originates with insurance companies, the vast majority of it does. Some of the very first types of insurance-linked securities were catastrophe bonds and catastrophe insurance derivatives. Their purpose, as is the purpose of most insurance-linked securities, is very simple: to transfer to the capital markets the risks that are too big for the balance sheets of insurance companies, or the risks that can be retained but whose transfer allows insurance companies to use their capital in the most efficient way. ILS such as reinsurance sidecars, value-in-force securities or securities designed to transfer excess reserves to the capital markets serve the same general purpose, with an emphasis more on capital management than on true risk transfer.

Insurance-linked securities serve as a link between the insurance industry and the capital markets. They provide insurance companies with new options in managing their risk and using their capital efficiently. Such direct transfer of insurance risk to the capital markets might not always be the best solution for insurance companies; however, it gives insurance companies another important tool that can be used in both risk management and capital management.

At the same time, most of the insurance industry has been unhappy with the development of such types of insurance-linked securities as life settlements and has seen this as “cannibalisation” of life insurance. Despite the initial negative reaction, it is likely that the industry will adjust to this development and might ultimately see it as a positive, since the transferability adds value to the life insurance product and can thus lead to growth in its sales. All insurance-linked securities make the markets more efficient, which is a positive for all parties.

INVESTORS

Investors never stop their search for yield. The search has intensified with the need to make up for the losses incurred during the 2008–2009 financial crisis and the realisation that traditional investment approaches are not going to accomplish this goal. The urgency of the search for sources of extra return is compounded by the growing emphasis on capital preservation and reduction in investment risk. These contradictory goals – maximising return and minimising risks – have always characterised the reality of investing. This duality has not changed, but the urgency of the first and the emphasis on the second have increased.

As unrealistic as it is, the desire to achieve high investment returns while taking low investment risks is as great as it has ever been. The Madoff affair demonstrated how very sophisticated investors might be willing to believe in the possibility of high returns delivered consistently, year after year, with very little volatility. People believe what they want to believe. The financial crisis of 2008–2009, however, brought fear to the markets, and the focus shifted from high returns to simple capital preservation. That fear remains, but we are now back to a situation where investors want high returns. The potential of high investment returns does exist, but in this quest there is a price to be paid in the form of greater risk. The choice of the right tradeoff between risk and return is as difficult as it has ever been.

In this environment, assets that have low correlation with the rest of the financial markets should be particularly attractive to an investor. Insurance-linked securities can serve the objective of capital preservation and contribute to portfolio diversification. While the common characterisation of insurance-linked securities as zero-beta assets is incorrect, many of them do have only weak correlation with the capital markets. The financial crisis demonstrated that for most types of ILS the relatively low degree of correlation with traditional financial assets stays low even under extreme circumstances, in the “tail” of the probability distribution where standard correlation assumptions tend to break down.

Insurance-linked securities, in particular those with a low degree of correlation with the financial markets, can be seen as a source of exotic beta. The exotic beta – the return associated with exposure to insurance as a risk factor only weakly correlated with the traditional markets – is really another form of alpha in the investment return. The ability to generate abnormal returns through this factor exposure should remain as long as the market inefficiencies exist. In insurance-linked securities, these inefficiencies are particularly great and likely to persist, in part due to the low level of investor expertise in the analysis of insurance risk. This situation makes insurance-linked securities all the more attractive to investors who currently do have the required expertise, as they can expect to generate sizable excess returns.

BOOK SCOPE AND STRUCTURE

This book has the simple objective of describing insurance-linked securities and insurance risk transfer from a practitioner's perspective – a viewpoint that is particularly important in a market that is new and still evolving. The book is designed to be a resource to those active in the marketplace, while also aiding basic understanding of the topics for those new to the field.

The scope was chosen to be very broad and to include all types of insurance-linked securities. While some hold the view that certain types of the securities described here do not belong in this category, choosing the broadest possible definition can only help in understanding the investment potential of ILS.

The book consists of five parts. Part I, "Introduction to Investing in Insurance Risk", provides an outline of the ways to obtain insurance exposure in investment portfolios. Insurance risk in general is discussed, after which "pure" insurance risk is defined and described. A brief overview of direct investment in insurance risk then follows, outlining the main types of insurance-linked securities. Motivation of both transferors and transferees of securitised insurance risk is also examined.

The next part, "Investing in and Modelling Securities Linked to Property and Casualty Risk", looks at the main types of securities used for transferring property and casualty insurance risk to the capital markets. It starts with an overview of cat bonds, which are the most widely known type of insurance-linked securities. Part II also describes derivative and derivative-type products linked to catastrophic events. An introduction to modelling catastrophe risk embedded in these securities is provided to help the investor to better understand their risk profile. Other types of insurance-linked securities, such as reinsurance sidecars and industry loss warranties,

are examined. A brief overview of weather derivatives is provided. Credit risk and other issues relevant to the analysis of property and casualty insurance-linked securities are also analysed.

Part III, “Securities Linked to Value-in-Force Monetisation and Funding Regulatory Reserves”, deals with insurance securitisations where the primary purpose is other than the transfer of insurance risk. Some such securitisations monetise the expected future cashflows from a book of insurance business, while others have to do with regulatory or accounting arbitrage. Not all of them fall under the strict definition of securitisation; in many cases, monetisation is the proper characterisation.

The following part, “Investing in and Modelling Securities Linked to Mortality and Longevity Risk”, describes securities that transfer to the capital markets the risk of mortality and longevity being different from expectations. Extreme mortality bonds, for example, are tied to the risk of a sharp spike in mortality. Derivatives linked to mortality risk are introduced, with a focus on catastrophe risk. These securities have a strong resemblance to the catastrophe bonds and catastrophe insurance derivatives described in Part II. Life settlements are discussed next, and it is explained how a life insurance policy can be viewed as a tradable asset. Some of the legal and accounting considerations involving life settlements are also introduced as they are particularly important for investors in life insurance policies. Key concepts in the modelling of mortality and longevity are outlined, with a focus on the issues relevant to analysing insurance-linked securities. Issues that have to do with longevity improvements and stochastic modelling of longevity are also described. Valuation of mortality-linked securities is discussed, with a focus on life settlements. Finally, longevity-linked securities are examined, with consideration of the role they can play in hedging the longevity risk of pension liabilities, life annuities, and portfolios of insurance-linked securities. While the primary emphasis is on longevity derivatives, other longevity-linked securities are discussed as well.

Part V, “Managing Portfolios of Insurance Risk”, deals with portfolio issues in the investment management of insurance risk. This final section reviews key aspects that have been touched upon in the preceding parts of the book, and describes a number of tools for managing securitised insurance risk on a portfolio basis. The first part of the section deals with catastrophe insurance risk. This is followed by a broader analysis of managing portfolios of insurance-linked securities of multiple types. Investment portfolio optimisation is discussed in the context of managing securitised insurance risk.

The Conclusion, following Part V, summarises the main ideas introduced in the book, and focuses on current trends in the insurance-linked securities market. It makes general observations about the market and discusses the expectations of how it will develop.

ACKNOWLEDGEMENTS

I would like to express my deep gratitude to all my friends and colleagues – too numerous to mention here – who have read all or parts of the manuscript, checked it for errors and provided helpful suggestions. Those contributing to the review process have included professionals from the asset management community, the insurance and reinsurance industry, investment banks, a modelling firm and an exchange, as well as brokers. The responsibility for any remaining errors or inaccuracies is mine.

Part I

Introduction to Investing in Insurance Risk

Investing in Insurance Risk

This introductory chapter provides a brief overview of concepts that are fairly obvious to most professionals in the investment and insurance field but may be unfamiliar to other readers. In addition, insurance risk is presented through an uncommon perspective that sheds light on its unique characteristics and corresponding investment considerations.

INVESTING IN RISK

There are no truly riskless assets. We always invest in risk. We might do it in the form of stocks, corporate bonds, real estate or treasuries, but ultimately the investment performance of these securities is predicated on their risks. We invest because we expect to earn a return commensurate with the risk we take in investing. In fact, we want the return to be higher than what the risk profile of an investment would imply.

Risk is good

It is too simplistic to say that “risk is bad”, and to think that it is something we want to avoid or minimise. Investing is always about risk. In fact, investors actively search for risk to invest in. As long as the compensation for taking on the risk is appropriate, the investment usually makes sense. A good investor is not the one who avoids risk; with excessive focus on avoiding risk such an investor will also strip out his return. A good investor is the one who invests in securities that together generate high risk-adjusted return appropriate to the investor’s goals. A good investor is certainly risk-averse, but only in the sense of not being willing to accept risk without proper compensation. As obvious as such statements may seem, the idea of seeking risk makes some uncomfortable. An investor must recognise that risk is good as long as it is the right kind of risk, the returns are commensurate with it and the overall investment objectives are satisfied.

The portfolio approach to investing is important to every investor. A pension fund might have allocations to individual asset classes and benefit

from the diversification it provides. The benefits of diversification explain why investments with low correlation to others are at a premium and being sought after. Certain types of insurance risk possess this desired quality of having low correlation with other asset classes.

INSURANCE RISK

Insurance risk lacks a clear, unambiguous definition. It is generally defined as the risk being taken on by insurance companies in selling insurance protection. This could be interpreted in a very broad sense to include all risks faced by an insurance company in the course of its operations. So it could be said that investing in insurance risk is the same as investing in an insurance company. Considered in this broad sense, insurance risk includes all traditional investment risks – market, credit, operational and others – as well as the insurance risk defined in a more narrow way – that of insurance claims (obligations under insurance policies) being greater than expected, or greater than a certain level that the insurance company wants or is permitted to take. Even this definition is imprecise, since all the risks are intermingled and cannot be fully decomposed into individual elements.

The more narrowly defined type of insurance risk would apply in cases of higher-than-anticipated losses due to factors such as random statistical fluctuations in the number of insurance claims or their severity, natural catastrophes or man-made disasters, spikes in mortality or fundamental shifts in longevity, and many others.

Often such types of insurance risk either cannot be transferred to investors purely through the traditional equity or debt instruments issued by insurance companies, or are best transferred to capital markets in a different fashion. Insurance-linked securities (ILS) are structured to transfer to investors this type of risk, and are specifically designed to address unique issues of insurance companies. Most have to do with the transfer of “pure” insurance risk where other risks are excluded or minimised. They afford investors exposure to risks that are different from those embedded in the traditional securities and that are often only weakly uncorrelated to the behaviour of the financial markets.

INSURANCE MARKETS

Before considering securities that are in some way linked or related to insurance, it is instructive to take a look at the insurance markets in general. Insurance markets have many unique features not found in other industries.

Insurance companies are highly leveraged enterprises in the sense that

their assets-to-shareholder equity ratio tends to be very high – particularly for life insurance companies and property-casualty companies in long-tail lines of business. While the degree of leverage across the capital markets has been going down (in some cases considerably), the insurance industry remains an exception.

The diverse offerings of insurance companies comprise two main categories: life and health insurance, and property-casualty insurance (referred to as general insurance in many parts of the world). Though these two categories of insurance are quite distinct, some companies handle both life and property-casualty insurance.

In describing insurance markets, it is also important to note that insurance is one of the most heavily regulated industries, a fact that, by itself, introduces a broad set of constraints and risks not found in other sectors. Moreover, the regulation to which insurance companies are subject is not uniform among jurisdictions, contributing to the fragmented nature of the insurance marketplace. Some jurisdictions impose price regulation and so insurance companies are not free to raise insurance rates on some of their products. In extreme cases, companies unable to raise rates for this reason have decided to exit certain products lines, but have encountered additional regulatory constraints, making this exit difficult. Few industries have to deal with such issues.

Another phenomenon specific to the insurance industry is the underwriting cycle; it pertains primarily to property-casualty insurance, and, to a lesser degree, to health insurance. Insurance companies as a group go through periods of charging customers rates that are too low, leading to rates of return dropping below the required level (referred to as “soft” markets); followed by periods when the companies are able to raise their rates to the level where they generate rates of return in excess of the minimum required (“hard” markets). This cycle does not have a simple logical explanation and is seen by many as evidence of how inefficient the insurance markets are. Arguably, no other sector has such a clearly pronounced profitability cycle, with the possible sad example of the airline industry. While many factors drive the underwriting cycle – changes in macroeconomic conditions, shock events resulting from investment losses or losses due to natural catastrophes, the fear of losing customer relationships, and many others – it is also recognised that some of the factors are purely psychological (such as the herd mentality). Predicting the next turn in the insurance underwriting cycle is a favourite pastime of the sell-side equity analysts who cover the insurance sector. The underwriting cycle

clearly is an important element in the analysis of most insurance-related investments.

Rating agencies play a crucial role in many types of insurance, as they assign insurance companies with financial-strength ratings, which are different from their counterparty credit ratings. The financial-strength rating reflects a company's claims-paying ability – that is, the level of certainty that policyholders will be paid when they make claims. Depending on the line of insurance, there are some critical thresholds below which an insurance company effectively finds itself unable to write new, or renew, insurance policies; falling even one notch below such a threshold can put a company out of business. This degree of vulnerability is not encountered in most other industries.

To sum up, insurance markets are unique because of a variety of factors, including fragmentation, particularly strict regulatory requirements, unusual risk and a significant degree of inefficiency. Deep understanding of such industry dynamics is a prerequisite to analysing many securities issued by this industry.

SECURITIES ISSUED BY INSURANCE COMPANIES

Insurance companies issue some of the same types of securities as do most companies in other industries. We can invest in insurance through common stock, debt or preferred stock. The analysis of the common stock of insurance companies is based on the general principles of equity analysis, while taking into account also the specific features of the insurance industry. Other types of securities issued by insurance companies are not found in most other sectors. An example would be surplus notes, which are securities similar to the trust-preferreds issued by banks. The securities issued by insurance companies are a relatively small part of the global capital markets, reaching at most 3% of their total size.

In the US, insurance companies provide two types of financial statements: traditional statements based on the Generally Accepted Accounting Principles (GAAP), and statutory statements mandated by insurance regulators. The volume of information contained in these statements is greater than what would typically be available for a company in another industry. Detailed exhibits provide a wealth of additional information. Both the GAAP and the statutory statements, along with other data released by insurance companies, help investors analyse the companies and value the securities they issue. The availability of the additional information, however, does not make the analysis easier and the uncertainty lower. There

are too many industry-specific issues that make the analysis different from that of other companies, and these issues present unique challenges. In the simplified analytical framework, we may often wonder why price-to-book ratios of insurance companies exhibit idiosyncratic behaviour, and what drives the difference in the price-to-book and other ratios between companies that appear to be rather similar based on their balance sheets, income statements and the business they conduct. Only a deeper level of analysis can answer such questions.


We may think that diversification can be achieved simply by investing in stocks or bonds issued by insurance companies, since they contain the “pure” insurance risk such as that of losses related to natural catastrophes or changes in mortality rates. However, these risks are rarely the main drivers of insurance stock performance. For most insurance companies, the main component of their profits stems not from underwriting income but from the investment returns on their asset portfolios. This explains why insurance companies, with their huge balance sheets and assets invested mostly in bonds and stocks, are heavily exposed to market risk. Life insurance stocks are seen by many as a beta play, as opposed to an uncorrelated asset.

Figure 1.1 overleaf illustrates the performance of the Dow Jones US Insurance Index relative to the S&P 500 Index and the Dow Jones Industrial Average. Correlation of the insurance index returns with the markets for the time period illustrated in Figure 1.1 was 80%, showing that investing in insurance stocks in and of itself does not necessarily provide diversification, because insurance companies are, to a significant degree, leveraged investment vehicles.

Warren Buffett puts it in slightly different terms by using the concept of float: “Float is money we hold but don’t own. In an insurance operation, float arises because premiums are received before losses are paid, an interval that sometimes extends over many years. During that time, the insurer invests the money.” This statement, repeated with minor variations in numerous annual letters by Buffett to the shareholders of Berkshire Hathaway, explains both the concept of leverage in insurance and why many insurance stocks have a high degree of correlation with the financial markets.

INSURANCE-LINKED SECURITIES

ILS are defined as financial instruments, other than traditional equity and debt securities issued by insurance companies, which carry insurance risk or a type of risk that is closely related to it. Examples of the risks included in insurance-linked securities are property-catastrophe risk, mortality,

Figure 1.1 Performance of insurance equities relative to stock markets

Source: Bloomberg

longevity and insurance loss reserve adequacy. ILS can also include many of the traditional risks such as market, credit and interest rate risks, but it is the inclusion of the significant degree of insurance risk that defines them.

The seemingly irrelevant question of what asset category ILS belong to is important. ILSs are normally classified as alternatives, but they come in many shapes and forms even for the same type of risk. These securities can be structured as fixed income instruments or as equities. Some ILS come in the form of derivatives while others most closely resemble private equity investments. A dedicated ILS fund can be limited to investing in only catastrophe bonds or have a broader mandate of investing in various types of insurance-linked securities and types of insurance risks they contain. The fund mandate determines how an investment in the fund itself is classified – whether it necessarily falls in the category of alternatives and, if the answer is positive, where it is placed within that category. The uncertainty as to the appropriate allocation bucket exists even in the cases of direct investment rather than that through a fund.

The classification may affect the flow of funds to ILS and insurance-linked strategies since they are relatively new and have not earned standard allocations afforded to the more traditional asset classes and investment strategies.

The size of the insurance-linked securities markets is very small relative

to that of the global financial markets, and even relative to the total value of securities issued by insurance-related entities. While exact figures are not available, the total size of the traded insurance risk (ILS), even when broadly defined and including both property-casualty- and mortality/longevity-linked securities, does not exceed US\$70 billion. Figure 1.2 shows estimates of the insurance-linked securities markets in relation to the broader financial markets.


Examples of ILS

The best-known example of an insurance-linked security is a catastrophe bond, or cat bond, with its primary investment risk linked to the occurrence of a natural disaster such as a hurricane or an earthquake. This security, like a corporate bond, pays coupons and returns principal to investors, unless it defaults. The default can be triggered by the occurrence of a natural catastrophe whose physical parameters have been specified in the bond covenants, or by insurance losses that exceed a predetermined level. In such cases, investors do not receive the full expected payments or any payments at all (which in some structures constitutes debt forgiveness). Cat bonds are structured in a way intended to minimise exposure to any risk other than the “pure” insurance risk of natural catastrophic events. In other words, every attempt is made to minimise the correlation with the conventional asset classes.

Another example of an ILS is embedded-value securities. Unlike a cat bond, which can be seen as a result of the securitisation of liabilities, embedded-value securities have to do with asset securitisation. Embedded-value securitisation is the exchange by an insurance company of its future profit stream on an existing book of insurance business for a monetary consideration received from investors now. In other words, it is a way to accelerate profits, which by itself is not unique to the insurance industry. These securities, contain a significant element of “pure” insurance risk but also include a number of other risks such as interest-rate and credit risks. While providing some diversification through exposure to insurance risk, embedded-value securities are certainly not zero-beta assets and have a significant degree of correlation with the markets.

INVESTING IN INSURANCE RISK

To summarise, securities issued by insurance companies can provide high returns but require specialised expertise for proper evaluation. Insurance markets have a number of unique features requiring adjustments to the

Figure 1.2 Relative size of the market

Notes: Derivatives, whose total notional amount is a multiple of the stock and bond markets combined, are not included. Estimates are as of 2009 and are based on data from the Bank for International Settlements, SIFMA, World Federation of Exchanges, World Bank, Milken Institute, World Economic Forum, LISA, Conning, and McKinsey. Only publicly traded securities are considered in estimating the size of the global financial markets. There are no adjustments for the cases of one public company owning stock of another publicly traded company. Securities that have connection to the insurance industry are broadly defined and include those issued by companies involved in other businesses in addition to insurance. A broad definition of insurance-linked securities is used to include such types of ILS as life settlements and industry loss warranties. Most private deals that can be reasonably characterised as ILS-type transactions are also included.

standard investment analysis used for most other industry sectors. The specialised expertise is a significant source of competitive advantage in the investment analysis of insurance equities and debt.

Investing in securities issued by insurance companies does not provide the diversification that might be expected from exposure to the risks of insurance losses being greater than expected due to the fluctuations in the frequency or severity of insurance claims, changes in mortality rates, or other risks unique to the insurance industry. In investing in corporate securities issued by insurance companies, most of the risks are not “pure” insurance risks but risks common to the financial markets. This explains the high degree of correlation between the investment performance of the insurance sector and the markets as a whole.

Search for uncorrelated return

Investors never stop their search for assets that improve the performance of their investment portfolios, either through extra yield or through exposure to uncorrelated assets. The value of truly low correlation with the markets became painfully obvious during the financial crisis that started in 2007. By

the end of 2008 all correlation assumptions broke down, and assets with historically low correlation all of a sudden started moving in sync. They were all moving in the same direction – down – and so were the supposedly diversified investment portfolios. Having investments with low beta generally improves portfolio risk-adjusted returns and contributes to the goal of capital preservation.

Insurance-linked securities as a portfolio diversifier

While insurance-linked securities are not zero-beta assets, they do represent a valuable and effective form of diversification. Many of them provide exposure to risks that have a low degree of correlation with the rest of the financial markets, while still generating a very competitive yield. Securities such as cat bonds issued after 2008, designed with an express intent to strip away, as much as possible, all risks besides the true insurance risk of natural catastrophes, provide a good illustration of this diversification.

A storm on Wall Street might shake the very foundation of financial markets, but it is not going to lead to a hurricane in Florida or an earthquake in California. A catastrophe bond is not going to be triggered because of the condition of the markets. The relatively low degree of correlation with market risk is the greatest advantage of insurance-linked securities, and for this reason insurance-linked securities can be an important component of most investment portfolios.

Insurance-Linked Securities

This chapter provides a brief introduction to insurance-linked securities (ILS) as an asset class in order to lay the foundation for the more thorough treatment of individual types of ILS in the rest of the book. It explains the reasons for transferring insurance risks to the capital markets and the benefits this transfer provides both to the insurance industry and to the investors. The types of insurance risk transferred to the capital markets are also briefly discussed.

INSURANCE-LINKED SECURITIES DEFINED

Chapter 1 defined insurance-linked securities as financial instruments, other than traditional equity and debt securities issued by insurance companies, that carry insurance risk or a type of risk that is closely related to it. Examples of the risks included in ILS are those associated with property catastrophe, mortality, longevity and insurance loss reserve adequacy. ILS can also include many of the traditional risks such as market, credit and interest-rate risks, but it is the inclusion of a significant degree of insurance risk that defines them.

The term “risk-linked securities” is occasionally used instead of ILS, sometimes to highlight a broader spectrum of insurance-linked securities – for example, weather derivatives, which do not have a direct relationship to any actual insurance losses, but serve the purpose of transferring to the capital markets risks very similar to those taken on by insurance companies. In some cases, the distinction between insurance-linked and other securities becomes blurred; but generally a security is labelled an ILS if it resembles one of the standard types of insurance-linked securities.

Insurance risks involved in insurance-linked securities cover the whole range of insurance-related risks, from property-casualty insurance to life insurance. The wide variety of insurance risks embedded in ILS is reflected in the multitude of types of insurance-linked securities.

TYPES OF INSURANCE-LINKED SECURITIES

While catastrophe bonds are the best known insurance-linked securities, the ILS universe is much broader than that. Products range from alternatives to reinsurance coverage, to securities that can be constructed only with the use of capital markets. Figure 2.1, opposite, presents ILS characterised by the degree of catastrophe risk being transferred to the capital markets and by the type of insurance risk. The list is far from complete: only the main types of insurance-linked securities are shown.

Categorisation of insurance-linked securities is partly dependent on the reasons the insurance risk is being transferred to the capital markets by insurance companies or other entities.

Reasons for transferring insurance risk to the capital markets

Insurance risk can be transferred to investors for a number of different reasons. Some of these reasons are described below. There is significant overlap since a transaction can accomplish more than one objective.

- ❑ **TRANSFER OF CATASTROPHE RISK.** Insurance and reinsurance companies are limited in the amount of true catastrophe risk they can assume. A large-scale catastrophe, either natural or manmade, has the potential of wiping out the surplus (shareholder equity) of many companies at the same time. It can even start a spiral of insolvencies or downgrades if several reinsurance companies fail, and the reinsurance recoverables remain uncollectable. Prudently managed insurance and reinsurance companies are aware of this risk and either partially transfer it to other parties or choose not to assume it at all, leaving some exposures uninsured. Since the total shareholder funds of the insurance industry are dwarfed by the size of the capital markets, it makes perfect sense to transfer the true catastrophe risk to investors. This can be done in the form of cat bonds, industry loss warranties, reinsurance sidecars, catastrophe derivatives, collateralised reinsurance of catastrophe risk, or contingent capital securities. Catastrophe risk also exists in life insurance – for example, in the case of a jump in mortality due to a pandemic event. Such risk can be transferred to the capital markets primarily in the form of cat mortality bonds and cat mortality derivatives.
- ❑ **SUBSTITUTE FOR TRADITIONAL REINSURANCE.** Limited risk capacity leads to higher reinsurance rates, which in some cases results in capital markets solutions being more efficient in terms of cost. Given the additional advantages provided by some insurance-linked securities (for example,

Figure 2.1 The broad range of insurance-linked securities and the insurance risks embedded in them

	Catastrophe risk	Non-catastrophe risk
Property and casualty insurance	Cat bonds	
	Industry loss warranties	
	Catastrophe derivatives	
	Collateralised reinsurance	
	Reinsurance sidecars	
	Contingent capital	
	Non-cat property and casualty bonds	
Life insurance and longevity	Cat mortality bonds	
	Value-in-force (embedded value) securities	
	XXX/AXXX securities	
	Life settlements and related securities	
	Cat mortality derivatives	
	Longevity bonds	
	Longevity derivatives	

the ability to lock in the cost of protection for more than one year, and limited credit risk), capital markets solutions can be an important part of the overall risk management programme, acting as both a substitute for and a complement to traditional reinsurance. Avoiding overexposure to a few reinsurers and thus lowering credit risk is of particular importance. Investor-provided collateralised reinsurance, insurance derivatives, and industry loss warranties are all examples of insurance-linked securities that fall in this category.

- ❑ **RELIEVING CAPITAL STRAIN.** In the absence of distressed conditions, insurance companies can still experience capital strain when they grow too fast or when regulations require them to hold capital significantly in excess of the levels necessary from the economic point of view. An example of a capital markets solution driven by this rationale is XXX and AXXX securitisation. In this case, US regulations require that reserves for some life insurance products be maintained at levels significantly in excess of what most consider economically reasonable. This requirement results in considerable strain on insurance companies' capital; XXX and AXXX securitisation or private investment solutions help alleviate this strain. Value-in-force securitisation or monetisation can also provide additional capital, either to eliminate a shortfall or to be used for other purposes such as mergers and acquisitions.
- ❑ **TURNING LIFE INSURANCE INTO TRADABLE INSTRUMENT.** Life settlements developed as a way for policyholders to monetise the value of their existing life insurance policies when they are no longer needed, when they cannot be afforded or when the benefit of immediate monetisation outweighs the advantages of keeping the policies. From the economic point of view, a life insurance policy is a security and thus can be traded. Once a life insurance policy is bought by investors, it then can be resold more than once. Portfolios of life settlements can be separately managed or securitised. Managing portfolios of life settlements can benefit from the use of another type of ILS, longevity derivative instruments, that could hedge the longevity risk of such portfolios.
- ❑ **LONG-TERM LONGEVITY RISK TRANSFER.** Capital markets solutions can be utilised to address the risk of greater-than-anticipated longevity. Pension funds and some annuity providers are among the entities exposed to this risk. In the case of pension funds, longevity improvements in excess of expectations can lead to significant shortfalls. Longevity derivatives and longevity bonds are examples of instruments that can transfer this risk to the capital markets.

The list illustrates some of the reasons why insurance risks would be transferred to the capital markets, along with a few types of insurance-linked securities used for this purpose. There are a number of additional reasons, including more efficient capital management, reducing earnings volatility of insurance companies, addressing rating agency concerns, managing credit risk and many others; again, these often overlap.

Reasons for investing in insurance-linked securities

While there is a multitude of reasons why insurance companies and other entities might want to transfer insurance risk, conceptually the reasons why investors might want to accept it are much simpler.

Adding an ILS to an investment portfolio may be beneficial if it improves the risk–return profile of the portfolio. Consequently, the analysis of whether an ILS investment makes sense is quite similar to the analysis of investing in any other security. If the marginal impact of adding an insurance-linked security to the portfolio improves its risk–return profile more than available alternatives, the investment probably makes sense.

In even simpler terms, investors find insurance-linked securities attractive because they provide yield, diversification or both. Given the constant search for extra yield and diversification opportunities, it is natural for investors to consider this asset class, with all of its unique characteristics. Structurers of insurance-linked securities are mindful of the fact that investor needs should be satisfied, and they take this into account when deciding on the best ILS structure to transfer an insurance risk to the capital markets.

YIELD AND DIVERSIFICATION OFFERED BY INSURANCE-LINKED SECURITIES

Investors look to insurance-linked securities primarily for yield or diversification. Diversification in particular has been publicised as a unique advantage of ILS. Insurance-linked securities do offer a type of diversification not available through exposure to other assets. For many types of ILS, especially cat bonds and similar instruments, this is a critical advantage that makes this asset class so important. The experience of 2008 shows that, when almost all asset classes are down, even those that historically have had low correlation, the importance of the low correlation that stays low even in the tail of the probability distribution becomes clearly evident.

The “zero-beta” assets

Many insurance-linked securities provide a unique type of diversification through exposure to “pure” insurance risk. While this is often their main attraction to investors, it does not mean they are completely uncorrelated with the rest of the financial markets.

Statements have been made repeatedly that ILS, in particular life settlements and cat bonds, are zero-beta assets and have no correlation with the markets at all. While the correlation between some types of ILS and the financial markets might be weak, it does exist, and the zero-beta claims are not valid. They are particularly unfounded where they are repeated most often – in the case of life settlements, which are clearly exposed to the interest rate and a host of other risks.

Yield generation

Insurance-linked securities often provide yield opportunities in excess of those implied by their risk level. The yield can be a very important benefit of these securities and can become an alpha generator for an investment portfolio.

Part of the reason for the extra yield is the market inefficiency and the unfamiliarity of investors with these securities. The market is still small, and expertise in ILS analysis is hard to find in the investment community. Over time, the markets will surely become more efficient, and excess returns will diminish or disappear. This, however, is likely to be a very long process.

Some ILS offering what appears to be high return on a risk-adjusted basis might in reality be much riskier than expected by investors lacking sufficient expertise in this space. Some of the ILS appear deceptively simple, and an investor without deep expertise in this asset class can be lured into making poor investment decisions.

Efficient frontier

The ability to invest in insurance-linked securities can have the effect of shifting the efficient frontier for an investor. The limited correlation of ILS returns with other assets enhances diversification options, and the new efficient frontier may then have lower risk for the same level of return, or higher return for the same level of risk. This is the exotic beta appeal of this asset class as it provides exposure to a risk factor with low correlation with the rest of the financial markets.

It is important that the efficient frontier mentioned above does not have to be defined within the mean-variance optimisation framework. In fact, the value of adding ILS to an investment portfolio can be even more apparent

in the more sophisticated framework that takes into account events in the tail of the probability distribution.

MARKET DYNAMICS

Despite its relatively small size to date, the ILS market is very dynamic and constantly changing. New instruments appear, or the existing ones suddenly grow in prominence, while others fade into obscurity, more or less in direct response to changing market conditions. Meanwhile there is a gradual, ongoing process of education and acceptance of this new asset class.

Not all of the developments have been smooth and the growth has been uneven. An example of such a change in the ILS market is the redesign of the cat bond structure to minimise the credit risk of this security. This was done in response to the realisation, driven by the events surrounding the bankruptcy of Lehman Brothers in 2008, that credit risk is present and can play a significant role in these securities. Another example is the uneven development of the life settlements market, which has been affected by problems specific to this asset class as well as by the general availability of risk capital in the changing investment environment. A further example is the painfully slow development of the longevity transfer market, despite the seemingly obvious need for it. Finally, exchange-traded catastrophe derivatives first appeared in the early 1990s but were unable to gain traction; now they have been reintroduced to address the growing needs of both hedgers and sellers of protection.

Demand for and supply of insurance-linked securities differ by the type of ILS and change over time, even for the same type of ILS. For example, reinsurance sidecars made a sudden appearance in the aftermath of the 2005 Katrina–Rita–Wilma hurricane season; they addressed an urgent need and then quietly decreased in importance. The existence of dedicated ILS funds brings another interesting element into the dynamics of this market, since they are effectively the source of captive capital that provides a guaranteed level of demand for some insurance-linked securities.

The financial crisis of 2007–2009 was a good test of the ILS market, as it allowed market participants to identify weaknesses of some of the ILS structures. More importantly, it underscored the general benefits of investing in most types of insurance-linked securities that provide both yield and diversification opportunities. It also drew attention to the need for proper expertise in the analysis of these financial instruments.

The convergence between the insurance and capital markets is occurring

slowly but steadily. Securitisation of insurance risk is an important part of this process. It addresses the needs of both the holders of insurance risk and the investors, and there is every expectation that the insurance-linked securities market will continue to grow and develop.

Part II

Investing in and Modelling Securities Linked to Property and Casualty Risk

Property Catastrophe Bonds

This chapter describes property catastrophe bonds, which are probably the best known type of insurance-linked securities. Standard structural features of catastrophe bonds are explained and the main analytical approaches introduced. The chapter explains advantages and disadvantages of these securities from both an insurance company and an investor perspective.

SECURITISATION OF PROPERTY INSURANCE RISK

The insurance industry is one of the largest warehouses of risk, incorporating the roles of both risk underwriter and risk bearer in the way that the banking industry did three decades ago. Since then, the banking industry has undergone dramatic changes and now passes much of the risk on to investors in the form of mortgage-backed and other securities. A strong argument could be made that the insurance industry should move in the same direction by underwriting insurance risk and then passing a sizable part of it on to investors in the form of standard securities. Many believe that this is eventually going to happen, in particular for the products that are more homogeneous and relatively commoditised, such as some types of life and automobile insurance. At this point, however, capital markets' involvement in the insurance industry is starting not from the standard homogeneous risk but rather from the most unusual and severe type of risk – that is, the risk of natural catastrophes.

Insurance and reinsurance industries, while considered to be well capitalised, do not have the capacity to withstand the financial impact of a large-scale natural disaster. Individual insurance companies, especially those with significant exposure in certain geographic locations, face the risk of large losses or financial ruin even from smaller-scale catastrophic events.

The sheer size of capital markets makes them the natural candidate for providing the backstop protection to the insurance industry should a Category 5 hurricane make a landfall in Miami, Florida, or should an earthquake Category 8 on the Richter scale hit San Francisco, California. Capital

markets, whose size exceeds that of the insurance industry by orders of magnitude, may more easily weather such catastrophic losses.

MOTIVATION FOR TRANSFERRING NATURAL CATASTROPHE RISK TO THE CAPITAL MARKETS

The idea behind catastrophe (cat) bonds is to transfer to the capital markets the risk that extreme catastrophic events would inflict sizable losses on portfolios of insurance policies held by insurance companies. Cat bonds offer a new way for insurance companies to manage their risk exposure, a way that provides benefits to insurance company shareholders by controlling the risk and, if used appropriately, deploying their capital more effectively. From the point of view of policyholders and regulators, the advantage is the decreased likelihood of the company's inability to pay its claims in the event of a natural catastrophe.

Insurance company motivation

The primary motivation of an insurance company in securitising its property catastrophe exposure by entering into a cat bond transaction is risk transfer. In contrast, in triple-X and most other life insurance securitisations, the primary motivation is not risk transfer but relieving the capital strain created by regulatory requirements. As part of the overall capital-management programme, the transfer of catastrophe risk to the capital markets is another tool that insurance companies have in their overall arsenal of ways to find the right balance between risk and return, and to manage capital more efficiently.

Cat bonds are used as an alternative to traditional reinsurance for low-probability events. In some cases, protection obtained this way is cheaper than the cost of reinsurance. An additional advantage is the fully collateralised nature of the cat bond protection. It reduces the credit risk that is always present in traditional reinsurance. This risk could be significant since, when a sizable natural disaster strikes, reinsurance companies are exposed to large losses and some might not be able to make good on their obligations. Cat bond transactions also allow insurance and reinsurance companies to lock in the cost of protection for a period longer than the one year that is standard for reinsurance contracts.

Investor motivation

The motivation of the insurer in hedging risk exposure is clear. What are the advantages of the transaction to the investor in these securities? In other words, why would capital markets players be interested in investing in cat bonds?

The first reason is the excess return that has been available on cat bond transactions. The excess (relative to similarly rated corporate debt) return has existed from the very first days of insurance risk securitisation and has been attributed primarily to market inefficiency. It has always been expected that with the growth of the cat bond issuance and the increase in the number and sophistication level of the market participants, the excess return would become very small. However, this has not happened in the decade since the first cat bond was issued even as we witnessed wide fluctuations in pricing. On the contrary, in the aftermath of the huge insurance losses in the 2004 and 2005 hurricane seasons, the excess return widened. This “Katrina effect” has led to investors’ being able to obtain high yields on securities that have relatively high credit ratings. The ubiquitous search for alpha has led some investors to this asset class.

The second and probably more important reason is the fact that cat bonds are often seen as almost “zero-beta” securities that provide a diversification benefit. The rationale behind this view is that cat bonds are weakly correlated with the other securities, leading to the comparison with Kipling’s “Cat That Walked by Himself”. For cat bonds that are properly structured, where all risks besides that of natural catastrophes are minimised, default rates are only slightly affected by movements in the financial markets. If the stock market crashes or the economy enters a recession, the effect on such cat bonds should be minimal. (In the past, most cat bonds included significantly greater credit risk than was intended by the structurers or appreciated by the investors. The bankruptcy of Lehman Brothers revealed this weakness in a painful way for some investors. The “new” cat bonds, issued since the beginning of 2009, have structural features than minimise the credit risk.)

HISTORICAL PERSPECTIVE

The idea of securitising insurance risk had been floating around for a long time before the first insurance-linked securities saw the light of day. Some of the first securities intended to transfer the insurance risk of natural catastrophes directly to investors were catastrophe options. Traded on the Chicago Board of Trade (CBOT) in the 1990s, they were met with lukewarm reception by both insurers and investors and were ultimately withdrawn. Exchange-traded catastrophe derivatives have recently reappeared and are now traded, in slightly different forms, on exchanges that include the Chicago Climate Futures Exchange, CME and Eurex. Chapter 5 provides more in-depth treatment of these securities.

Cat bonds have enjoyed greater success. One of the first cat bonds was

issued on behalf of USAA, a large insurance company, in 1997. It transferred to investors the risk that a hurricane in the Eastern US and the Gulf Coast would result in catastrophic insured losses to the company. The size of the bond was US\$395 million, which was the maximum protection size provided to USAA by the transaction.

Since that pioneering transaction, the volume of property catastrophe insurance securitisations has steadily grown, primarily in the form of catastrophe bonds. Insurance and reinsurance companies as well as corporate entities have turned to the capital markets for protection against catastrophe risk. The type of risk transferred to investors has ranged from hurricanes to earthquakes to typhoons, in geographic areas spanning the globe from the US to Europe to Japan.

Until recently, the growth was not as fast as observers had anticipated. However, in the aftermath of Hurricane Katrina in 2005, the interest in securitising property catastrophe insurance risk has exploded, and there has been dramatic growth in the total amount of capital invested in securitised risk in the forms of cat bonds, industry loss warranties and reinsurance sidecars. The last two are described in greater detail in Chapter 6. The temporary pause in issuance in the second half of 2008 had to do with the above-mentioned credit risk issues, which have now been largely resolved.

RISK TRANSFER IN INSURANCE

Insurance loss distributions tend to differ significantly from the normal distribution (the bell curve). They are referred to as fat-tailed distributions because of the high probability of extreme diversion from the mean. (More precisely, these are leptokurtic distributions. Their excess kurtosis leads to the higher probability of outliers in a sample relative to samples drawn from a Gaussian distribution.) Insurance losses resulting from natural catastrophe events lie at the far-right tail of the aggregate loss distribution, the “cat’s tail”. These events and their financial impact are difficult to model but are important for insurance companies to protect against.

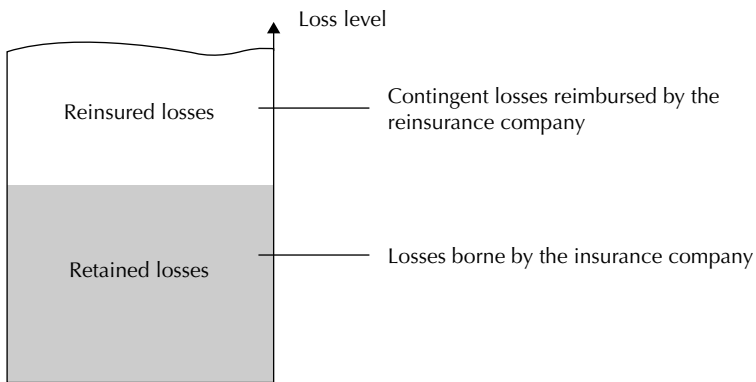
Insurance companies often find themselves unable or unwilling to retain all of the risk inherent in their portfolios of insurance policies. In dealing with catastrophe risk, the two main mechanisms for risk transfer are reinsurance and, more recently, cat bonds or similar capital markets solutions. Reinsurance plays a very important role by providing a somewhat efficient risk exchange mechanism for the insurance industry. In dealing with large-scale catastrophic events, however, even reinsurance fails to provide

adequate protection due to the limited capital in the reinsurance and insurance industry relative to the magnitude of potential losses.

Reinsurance risk transfer


Discussion of risk transfer and catastrophe bonds is impossible without describing reinsurance, the main mechanism for risk transfer in the insurance industry. Simply put, reinsurance is insurance for insurance companies. In the case of catastrophe risk transfer, an insurance company can buy reinsurance protection against losses exceeding a certain level.

Figure 3.1 Simplified example of a catastrophe reinsurance structure

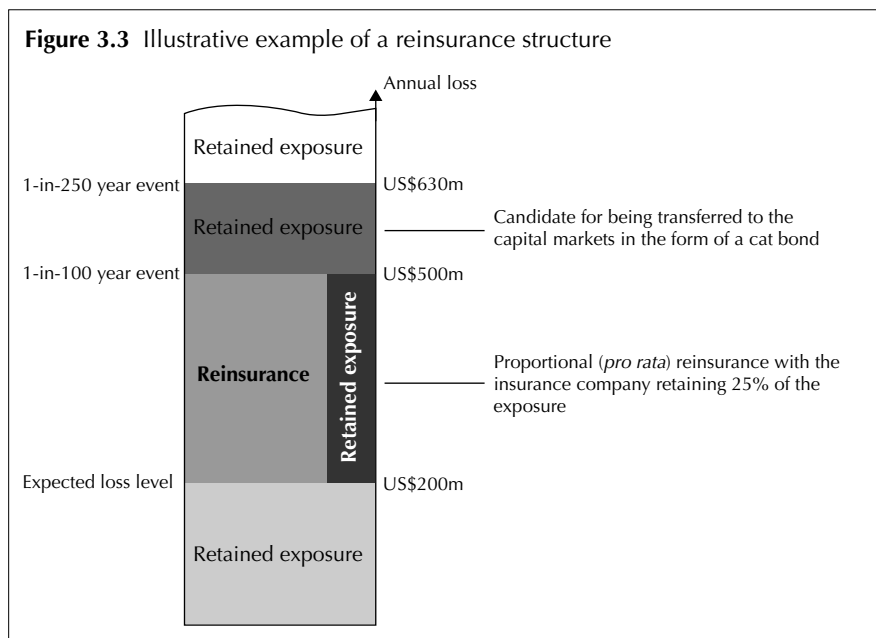


The insurance company, or “cedent” of risk in the reinsurance parlance, pays premiums to a reinsurer for the protection, and is reimbursed for claims in the scope of the reinsurance contract.

Figure 3.2 Reinsurance cashflows



Depending on jurisdiction, reinsurer’s rating and other considerations, the reinsurance company might be required to post collateral. Reinsurance companies could in turn reinsure some of their risk. This type of reinsurance is called retrocession.



In the example presented in Figure 3.3, if the company decides that it is too risky to retain the exposure between US\$500 million and US\$630 million, it has two main options. One of them is to reinsure this exposure. Another is to go to the capital markets and obtain protection in the form of a catastrophe bond or a similar instrument. In addition, there is always an option to reduce the insurance risk exposure by either writing less insurance business or by changing the concentrations, policy limits or policy conditions of the insurance portfolio. There are also options of raising additional capital in the form of equity, debt or hybrid securities, as well as obtaining contingent capital. From the point of view of the efficient use of capital, these options are usually less effective than reinsurance or catastrophe bonds. Advantages and disadvantages of using cat options and futures to protect against catastrophic events are discussed in Chapter 5.

CATASTROPHE BOND STRUCTURE

The structure of a cat bond is different from that of asset-backed securities. Effectively, securitising insurance risk amounts to securitising a liability rather than an asset.

Unlike the case of corporate bonds, the insurance or reinsurance company transferring catastrophe risk to the capital markets is not issuing the bond

directly. Instead, the bond is issued by a special purpose reinsurance company, which is generally located offshore. Thus, the entity that transfers the risk to the capital markets is referred to as the sponsor rather than the issuer of the catastrophe bond.

An entity that wants to transfer catastrophe risk to the capital markets would enter into a catastrophe reinsurance contract with a special purpose vehicle (SPV), a reinsurance company. The SPV will issue a bond with the payment of principal and interest contingent on there not occurring a catastrophe causing specified damage. The term of the reinsurance contract is the same as the term of the bond. If during this term no such catastrophe has happened, investors get back the principal and interest in full. Should there be a natural catastrophe triggering the reinsurance contract, the SPV will pay the claims. The remainder of the funds, if any, will go towards the payment of principal and interest to investors.

The simplified structure of a catastrophe bond is shown in Figure 3.4.

If no covered catastrophe has occurred during the term of the bond, investors receive back their principal at the end of the term.

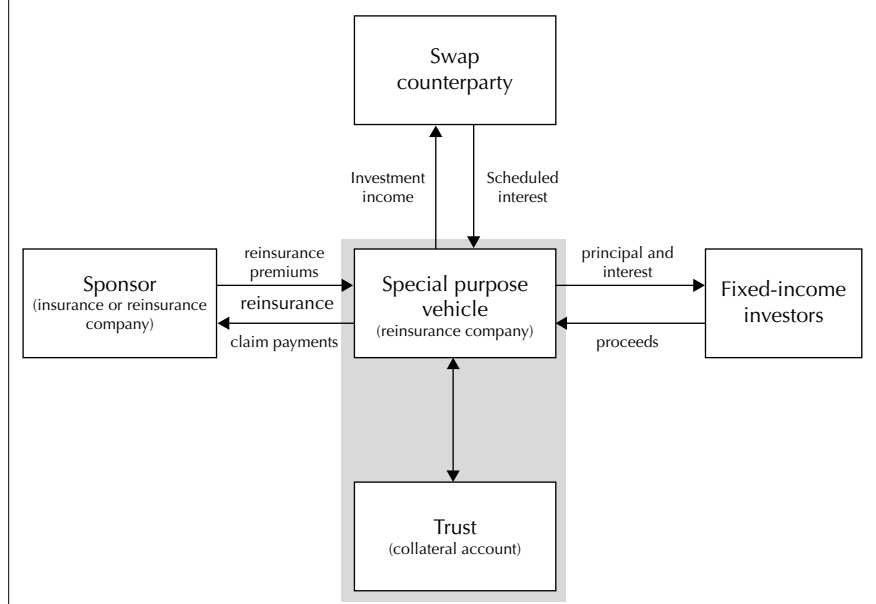
The structure involving an SPV is commonly referred to as “legal separation”. The SPV issuing the bond is a bankruptcy-remote entity. This allows the cat bond to be issued on a non-recourse basis. The insurance company sponsoring the bond does not have a claim on the assets; if the insurance company goes bankrupt, the investors are not negatively affected.

In another structure, the insurance company issues the bonds directly without using the SPV mechanism. Such a bond would generally be subject to recourse, putting investors at risk should the issuer suffer an insolvency. The legal separation structure is preferred by investors and has become standard in cat bond issuance.

In a more detailed cat bond structure (see Figure 3.5), an SPV (a reinsurance company), simultaneously enters into two transactions. The first is a reinsurance contract with the sponsoring company calling for the reimbursement of insurance losses above a specified level, with the losses caused

Figure 3.4 Simplified structure of the catastrophe bond flow of funds



Figure 3.5 Typical structure of a catastrophe bond

by a certain natural catastrophe. The second transaction is the issuance of a fixed-income security, a cat bond, to investors. The cat bond provides for payment of interest and repayment of principal unless a default is triggered by a natural catastrophe leading to a high level of insured losses.

Proceeds from the sale of the cat bond are deposited into a trust account that serves as collateral. The trust account would contain very secure, highly rated short-term instruments. While in many cases a cat bond sponsor could legally own the SPV without affecting its bankruptcy-remote status, in practice the SPV would usually be established by a third party such as an investment bank structuring the transaction.

Returns from the collateral account are swapped for a Libor-based rate with a highly rated counterparty. The total-return swap feature has become common in cat bond structures. Thus, interest rate-risk is minimised and the cat bonds become floating-rate instruments.

The interest payments received by investors are composed of the Libor-based returns on the funds in the collateral account; they can also include all or part of the reinsurance premiums received by the SPV from the sponsor.

Several ways to minimise credit risk related to the swap counterparty and to the assets in the collateral account have emerged post-2008. They are

described in Chapter 7, along with modifications to the cat bond structure that accomplish this goal.

No credit enhancement or credit wrapping has been used in property catastrophe bond securities. This has to do, in part, with the relatively low ratings of most catastrophe bonds, which makes them too risky for mono-line financial guarantee companies to add a credit wrap. (Credit enhancement used to be a common feature of extreme-mortality catastrophe bonds described in Chapter 11 and other life-insurance-linked securities. Credit enhancement of this type is no longer available from financial guarantors. See Chapter 7 for discussion of credit risk in cat bonds and other insurance-linked securities.)

DEFAULT TRIGGERS

A number of payout triggers – triggers of the cat bond default – have been proposed and used in cat bond transactions. In general, the triggers fall into one of two categories: indemnity and index.

Indemnity triggers

Indemnity triggers provide for cat bond payout based on the actual insurance losses suffered by the bond sponsor. This makes the cat bond a very effective hedge against the risk of losses from the natural catastrophe since the basis risk is minimised. It largely avoids the unfortunate situation of a natural catastrophe occurring, an insurance company suffering significant losses, but finding itself unable to collect from the cat bond it sponsored.

The negative side of indemnity triggers, from the point of view of the sponsor, is the need for information disclosure about its book of insurance business and underwriting practices. Many insurance companies prefer to keep this data confidential. Some of them are also hesitant to undergo the data quality review needed to present the information to investors in an offering circular.

Many investors see only negatives in the use of indemnity triggers. By its very nature, an indemnity trigger is less objective since it is based on actual insured losses rather than on parameters of a physical event. Investors are justifiably wary of the asymmetric information, with the insurance company sponsoring the bond having a significant information advantage in better knowing the types of risks it underwrites, risk aggregation, its underwriting standards and claim-settlement practices. The investors also assume the risk that, as the company implements its strategy or responds to market conditions during the term of the bond, its insurance portfolio might change and

increase the risk of bond default. The very fact that the bond sponsor has obtained protection via a cat bond could lead to a morale hazard, demonstrating itself in less care being taken in insurance underwriting and claim settlement. In addition to the morale hazard, there is always a potential for moral hazard, with the insurer intentionally (but without violating the bond covenants) making changes to its portfolio to the detriment of the cat bond investors.

While indemnity-based bonds historically were the first issued and are still common, the general trend has been away from the indemnity-type triggers and towards index triggers.

Index triggers

Index triggers do not directly depend on the bond sponsor's actual insurance losses. Rather, they depend on parameters that are outside of the control of the sponsor, thus providing more comfort to investors by eliminating the information asymmetry inherent in indemnity-based triggers. Index triggers usually fall into one of the following four categories: simple index, parametric, model portfolio loss and industry loss.

Basic index trigger

Basic index trigger provides for cat bond payout in case a predetermined physical event happens. A simple example would be a Category 5 hurricane making a landfall in Florida. If such an event happens, a cat bond with this trigger will suffer a default and make a payment to the benefit of the sponsor. It could, but does not have to, be structured like a binary option, providing either no payment to the sponsor if not triggered or the full payment (full default) if triggered.

From the point of view of an investor, this structure is very attractive. Investors have access to full information, and the dependence on sponsor's underwriting and other practices is eliminated.

On the other hand, the insurance company sponsoring the bond faces significant basis risk. A trigger so crudely defined could have poor correlation with actual insurance losses, reducing the effectiveness of the cat bond as a hedge. In other words, there is a significant chance that the cat bond would provide little or no protection against actual insurance losses suffered by the sponsor. There is also a chance that the bond will be triggered when the sponsor has not suffered sizable losses. In this case, the sponsor has paid for unneeded protection.

The basis risk is present in all non-indemnity trigger types, but is greatest when the trigger is based on a basic index.

Parametric trigger

Parametric trigger is based on the occurrence of catastrophic events with a combination of defined physical parameters. More than one type of catastrophic event (hazard) could be involved, and the amount of the payout is a function of the physical parameters of the cat events. A predefined formula is used to determine whether the bond is triggered and what the payout amount is. The formula could be quite complex. It is structured in a way that reduces the basis risk by identifying physical parameters of cat events (such as wind speeds at several locations) that would lead to insurance losses of the magnitude that the sponsor wants to transfer to capital markets. Identification of such parameters and the construction of the formula (the overall index), if done properly, involve a significant modelling exercise on the part of the sponsor. The investor, on the other hand, is not concerned with the sponsor's insurance losses and hedge effectiveness. Since the probability of default and the loss given default are independent of actual insured losses, investor analysis is focused on the probabilities of the physical events and their severities included in the parametric trigger formula.

Model portfolio loss

In this case, a sponsor creates a model portfolio that closely mirrors its actual portfolio of insurance policies or the portfolio that the sponsor expects to hold during the term of the bond. The portfolio is held "in escrow" together with the modelling software used to calculate losses to the portfolio. If a natural catastrophe happens, its actual physical parameters are input into the modelling software and losses to the model portfolio are generated. The bond payout depends on whether and by how much the modelled losses exceed a predetermined level.

To further reduce basis risk, the sponsor could use its actual current insurance portfolio instead of the representative model portfolio. The negatives of this approach have to do with the unwillingness of insurance companies to reveal detailed information about their insurance portfolios and the fact that such detailed policy-level disclosure could sometimes be unlawful.

Investors not possessing specialised expertise and knowledge of cat modelling software sometimes feel uncomfortable with the use of this trigger type.

Industry loss trigger

This trigger is tied to an index of losses suffered by the insurance industry as a whole as a result of a natural catastrophe. While not based directly on

physical parameters of a catastrophic event, this index could be modelled much better than indemnity losses. Given that a specific catastrophic event has occurred, insurance losses for the whole industry are more predictable than losses for an individual insurance company. They also are not subject to manipulation by the sponsor through claim settlement or another mechanism. Insurance loss-reporting organisations provide information to determine the overall loss level for the industry from a specific catastrophic event. The sponsor bears the basis risk, which depends on how its actual loss distribution differs from the rest of the insurance industry.

Trigger choice

In choosing a trigger, there is always a balance to be struck between transparency and simplicity on the one side, and the minimisation of basis risk on the other.

It is also worth noting that trigger choice to some degree affects structuring costs, with indemnity-based transactions being the most expensive to structure. Indemnity-based cat bonds also take longer to generate a payout since the sponsor might have to settle its claims first to determine the loss size. Basic index, parametric and model portfolio triggers provide for fast payout, while cat bonds based on industry-loss triggers have a payment delay due to the need to calculate the estimates of industry losses.

While in general all default triggers fall into one of the described categories, modifications of these triggers could be and have been used too.

Some investors, especially in the aftermath of 2005 Hurricane Katrina, have expressed a strong aversion to indemnity-based transactions, and prefer bonds with parametric and similar triggers.

Second- or third-event trigger

Structuring a cat bond provides a lot of room for creativity in trying to achieve the best protection for the insurance company while satisfying investor concerns. Sometimes an insurance company is not afraid of suffering one catastrophic loss, but wants to get protection in case one catastrophe is followed by another in the same or the following year. A second-event trigger could provide the required protection to the company, with the bond providing no payout (but being “activated”) after the first catastrophic event and paying only if the second event occurs as well.

NUMBER AND TYPES OF PERILS

A catastrophe bond trigger could be based on one specific type of peril such as a hurricane, typhoon or earthquake. It could also be based on a number

of perils, with losses from any one of them or a combination of perils triggering the payout.

While the first cat bonds were generally designed to provide protection against one type of peril, we have now seen a strong trend towards incorporating multiple perils in a bond. The same bond could have a number of peril/geographic location combinations. Larger insurance or reinsurance companies with portfolios of insurance policies on more than one continent are interested in this aggregate protection. When it comes to multiple perils, investors fall into two categories. Some are happy to see various types of uncorrelated risk in the same security. Effectively, diversification is provided for them in such a bond. Others prefer to buy cat bonds tied to a single peril and to achieve diversification on their own. The latter category tends to include investors with better understanding of the insurance-linked securities, including the funds focused exclusively on these financial instruments. Table 3.1, overleaf, shows a sample of catastrophe bonds issued based on various default triggers and types of catastrophe peril. Some cat bonds have included a number of tranches, each of which corresponds to a specific type of insurance risk and has its own trigger.

The securitisation of insurance risk has moved beyond property catastrophe and has included some liability insurance cat bonds, as well as securitisation of property-casualty insurance risk that is not truly catastrophic in nature. (Securitisation of extreme mortality risk is discussed in Chapter 11.)

TERM

The cat bond tenor varied widely in the early days of insurance securitisations, but has now stabilised with the average being three years. This term is long enough for the sponsor to lock into a multi-year protection at a predetermined price and to avoid paying the fixed cost of issuing a cat bond every year. At the same time, it is short enough for the sponsor to predict the composition of its future insurance portfolio with a reasonable degree of confidence.

QUANTITATIVE ANALYSIS

Both investors and the sponsor require a good understanding of potential losses, that is, the probability distribution of cat bond payouts. This probability distribution is in turn based on probabilities of the cat bond being triggered, and the payout amounts given that the bond has been triggered.

Table 3.1 Representative catastrophe bond transactions for various default triggers and types of peril

Cat bond sponsor	SPV	Year	Type of peril	Type of trigger
USAA	Residential Re I	1997	Hurricane in Eastern/Gulf States	Indemnity
Tokyo Marine and Fire	Parametric Re	1997	Earthquake in Japan	Parametric
Vivendi Universal	Studio Re	2002	Earthquake in California	Modified parametric
Oil Casualty Insurance	Avalon Re	2005	Industrial accident (excess liability)	Indemnity
Swiss Re ¹	Kamp Re	2005	Hurricane and other (multi-peril)	Indemnity
Swiss Re	Successor Class B	2006	US windstorm	Modified model portfolio loss
SCOR	Atlas Re III	2006	European windstorm and earthquake in Japan	Second and subsequent event
Endurance	Shackleton Re Class A	2006	Earthquake in California	Industry loss
Allianz	Blue Wings	2007	UK flood, and earthquake in US and Canada	Combination of modelled loss and parametric
State Farm	Merna Re	2007	Earthquake, hurricane and other in US and Canada	Indemnity aggregate
Flagstone Re	Valais	2008	Combination of natural cat perils in several countries	Indemnity
Allstate	Willow Re	2008	Texas hurricane	Modified industry loss
SCOR	Atlas Re V	2009	US and Caribbean earthquake and hurricane	Modified industry loss
Travelers	Longpoint Re II	2009	US hurricane	Modified industry loss
Hartford	Foundation Re III	2010	US hurricane	Modified industry loss


Exceedance curve

Insurance-linked securities might be the only asset type for which probabilistic risk analysis is included in the investor prospectus. For catastrophe bonds the analysis, performed by one of the firms specialising in modelling catastrophe events and their financial impact on portfolios of insurance policies, is usually presented in the form of a probability exceedance curve (EC). The exceedance curve shows probabilities of insurance losses of various magnitudes.

If the annual exceedance probability is 1%, then the probability of exceedance during a three-year period is 3%. (More precisely, the probability of exceedance over a three-year period is equal to $1 - (1 - 0.01)^3 = 2.97\%$. The approximation works well for only very small annual exceedance probabilities and short time periods. For example, if the annual exceedance probability is 2% and the term is eight years, we might think that the probability of exceedance over the term equals 16%. In reality, it is 14.92%, which is calculated as $1 - 0.98^8$.) Figure 3.6 shows an example of an exceedance probability curve for a portfolio of insurance risk.

In this example, losses above US\$500 million might have a catastrophic effect on the insurance company's financial position. The company has several options to protect itself against this possibility. Some of them have to do with raising additional capital or reducing or rearranging the company's portfolio of insurance policies. The most common solution is purchasing reinsurance – that is, insurance protection for this insurance risk portfolio. For example, the reinsurance coverage could take the form of the reinsurance company reimbursing the insurance company for all losses above the level of US\$500 million, limited to the total payout of US\$250 million. In this

Figure 3.6 Exceedance probability curve for a portfolio of insurance risk




case, the insurance company would still be unprotected if the total losses exceed US\$750 million, but would probably be willing to take this risk if losses above US\$750 million were considered to be exceptionally unlikely. The company might wish to protect itself from losses in excess of US\$500 million even if the effect of such losses would not have a truly catastrophic effect on its financial position. The reasons for it could be the desire to decrease earnings volatility or to reduce capital requirements.

In property insurance, unique terminology has been developed. Probable maximum loss, or PML, is the loss level that would be reached only extremely rarely. There are many opinions of how rare is “rare”, leading to multiple definitions of PML. If a company wants to define PML as the aggregate loss level that would be reached only once in 100 years, then in the example above the PML will be US\$500 million. With the increased emphasis on risk management and the more stringent capital-adequacy requirements promulgated by the rating agencies, there is growing shift of focus to property catastrophe events that happen on average less often than once in 250 years, leading many to define PML as the 1-in-250-year event. While the concept of PML is often used in relation to losses from individual policies, here we discuss the aggregate PML of an insurance portfolio. We also avoid non-quantitative definitions of PML still common in the insurance industry.

In insurance, probability of exceedance is usually expressed on an annual basis, that is, as a probability that insured losses will exceed a certain level

Figure 3.7 Exceedance probability and probable maximum loss for a portfolio of insurance risks



over a period of 12 months. In the context of cat bonds, exceedance probability could also be expressed as the probability of losses exceeding a certain level, such as the bond trigger level, over the term of the bond.

Chapter 4 provides more details on modelling catastrophe risk. Modelling catastrophe risk presents numerous challenges, but, even when it is accomplished, the results by themselves do not tell the investors what price is appropriate or fair for the cat bond being modelled. Several pricing models have been proposed. Often, they use as an input the observed prices for other cat bonds. An example of a pricing model is the Wang transform introduced in Panel 3.1, overleaf. (This panel is meant to introduce the concept of the Wang transform; its full explanation is outside the scope of this chapter. As with all panels in the book, it can be skipped by the reader without detriment to the overall understanding.)

As “neat” mathematically as the Wang transform is, its practical application is very difficult. It has also been pointed out (Pelsser 2008) that its use in pricing financial and insurance risks is consistent with arbitrage-free pricing, only under rather restrictive assumptions (this statement, however, has been disputed).

Other pricing approaches have been proposed, such as the application of extreme-value theory to cat bond pricing. This approach requires making assumptions not fully appropriate for cat bond analysis, and it does not produce results resembling observed cat bond prices. A simple rule-of-thumb approach to pricing includes the use of “multiples” of expected annual loss (average annual loss, or AAE) to determine the required spread over Libor or risk-free rate. Different multiples correspond to different levels of expected loss. While this approach has a questionable mathematical foundation, it is easy to use and there are some investors that utilise it. Another simple approach that has been proposed calculates prices based on the expected frequency and severity of the losses. The parameters are estimated based on the observed cat bond prices. This approach has the appeal of simplicity, but it lacks any theoretical foundation. Finally, some still use approaches that calculate prices based on the mean plus a multiple of standard deviation. Many of these relatively simple approaches are borrowed from reinsurance pricing, where they have been used for many years, but even there they are being replaced by the more sophisticated methods.

In addition to the shaky theoretical foundations of some of the pricing approaches, their common weakness is the dependence – either for parameter fitting or for results validation – on the actual observed cat bond prices. The cat bond market and the ILS markets in general are far from

PANEL 3.1 WANG TRANSFORM AND PRICING OF CAT BONDS

The Wang transform was developed by Shaun Wang (Wang 2000; Wang 2004; Pelsser 2008) with the goal of linking actuarial pricing and modern finance theories. It has been used for both pricing cat bonds and excess-of-loss reinsurance. While the full explanation of this method is outside the scope of this chapter, the basics of the approach are explained below.

Based on the underlying loss variable X , the loss to the excess-of-loss layer attaching at a with the limit of h , which is equivalent to the loss to a cat bond, is defined as

$$X_{[a, a+h]} = \begin{cases} 0, & \text{for } X < a \\ X - a, & \text{for } a \leq X \leq a + h \\ h, & \text{for } a + h \leq X \end{cases}$$

For a loss exceedance probability $S(x)$ over the interval $[a, a+h]$, the expected loss is

$$E[X_{[a, a+h]}] = \int_a^{a+h} S(x) dx$$

For a very narrow layer (very small limit h) this can be written as

$$E[X_{[a, a+h]}] = S(a)h$$

The price for this layer, $E^*[X_{[a, a+h]}]$ contains, in addition to the expected layer loss $E[X_{[a, a+h]}]$, a risk load. Price-based (or risk-adjusted) loss exceedance probability is then defined by Wang as

$$S^*(a) = E^*[X_{[a, a+h]}] / h$$

Wang proposed the following transform to obtain $S^*(x)$ from the loss exceedance probability $S(x)$:

$$S^*(x) = \Phi[\Phi^{-1}(S(x)) + \lambda]$$

Here Φ is the standard normal cumulative distribution, and λ is a parameter closely related to the Sharpe Ratio. Treating liabilities as negative assets, the Wang transform for the asset gain viable X then becomes

$$F^*(x) = \Phi[\Phi^{-1}(F(x)) + \lambda]$$

where $F(x) = 1 - S(x)$ denotes the cumulative distribution function of X . The exceedance probability distribution can have any form in this formulation. Based on the normality assumption for $S(x)$, λ is equal to the Sharpe ratio. In catastrophe risk, the distribution is not normal and can be very skewed and have excess kurtosis. Using observed prices, we can utilise the Wang

transform for pricing cat bonds and cat risk in general. Wang further introduced a technique to modify the transform, to account for the very fat tails of the distribution by incorporating additional risk adjustments, and to reflect the risk premium appropriate for higher moments of the distribution.

being efficient, and the observed prices, even relative to each other, do not necessarily follow the logic evident in more efficient markets.

The supply–demand dynamics play a very important role in pricing cat bonds and catastrophe risk in general. When reinsurance markets “harden”, the spread over Libor is likely to increase. This effect does not necessarily correlate with the behaviour of the financial markets. Even more importantly, the “peak peril” effect results in prices that are difficult to predict based on the assumption that markets are efficient. Two cat bonds, one linked to hurricane losses in Florida and the other to typhoon losses in Australia, might have exactly the same exceedance probability distributions, but the yield on the Florida hurricane bond is likely to be dramatically greater than on the Australia typhoon one.

While the proposed pricing approaches often fail in the analysis of individual bonds, relative-value analysis is still possible and helpful. More importantly, the existing modelling tools allow us to manage cat risks on a portfolio basis, and – instead of trying to come up with a theoretically correct price for an individual bond – to see what incremental impact its addition to the portfolio is going to have relative to the available alternatives. This topic is further discussed in Chapter 16.

Return period

Often, the data is presented in the form of return period instead of exceedance probability. These two terms are closely related. Return period is the average length of time between occurrences of events exceeding a specified threshold. If the annual exceedance probability is 1%, the return period is 100 years.

As with the probability exceedance curve, we can draw a graph of return period as a function of loss level, and base decisions on the data presented in this format.

Stress testing and sensitivity analysis

While the quantitative analysis is based almost entirely on the probability exceedance curves produced by catastrophe modelling software, scenario

testing is often utilised too. It is used in part as a check on the “black box” software used to model catastrophe insurance losses, and in part as a stress-testing mechanism. For example, for an insurance risk concentrated in Northern California, one might want to estimate the losses that would be incurred if the 1906 San Francisco earthquake happened today.

Stress testing, by necessity, has to be performed using the same modelling tools as those used to produce the probability distribution of catastrophe losses. Since no other tool is available, stress testing often involves moving along the probability curve and evaluating the results of a catastrophic event that the model considers less likely.

Sensitivity analysis could be performed in the standard way of varying the input parameters of the model and observing the effect on the probability exceedance curve and losses affecting the cat bond. Ideally, more than one type of catastrophe modelling software would be used to produce probabilistic results that could then be compared. While it is sometimes done by the sponsor of a cat bond, this data rarely finds its way to investors. The so-called cat bond remodelling process introduced by the three major catastrophe-modelling firms attempts to alleviate this informational deficiency; and is described in Chapters 4 and 16.


INVESTMENT PERFORMANCE OF CAT BONDS

Ever since the first cat bonds, insurance-linked securities have been issued at widely fluctuating yields. Such market inefficiency is normal for any new type of security, in particular if the market is still developing and lacking real liquidity. As a group, catastrophe bonds have outperformed many other securities bearing the same degree of risk, when risk is defined only in terms of probability of default and loss in the case of default. (In the cat bond vernacular, these are called “attachment probability” and “conditional expected loss”.) More importantly, their volatility has been lower and correlation with the markets weaker than for most other fixed-income securities. This stellar performance, however, suffered in 2008, when there emerged credit-risk issues in cat bonds (though these were corrected in the newer structures described in Chapter 7, and when the forced selling by multi-strategy hedge funds temporarily depressed cat bond prices in the secondary market.

Historical performance

Figure 3.8, overleaf, shows investment weekly performance of publicly disclosed catastrophe bonds relative to the corporate debt with the same ratings.

Figure 3.8 Investment performance of cat bonds relative to other fixed income securities



Source: Bloomberg (Swiss Re Cat Bond Total Return Index SRCATTRR, Barclays Capital US Aggregate Index (investment grade), Barclays Capital US Corporate High Yield Index)

Excess spread

Spreads for catastrophe bonds have historically exceeded those for comparably rated corporate securities. There are multiple reasons for the extra spreads enjoyed by cat bond investors. The most important of these are the following.

- ❑ **NOVELTY PREMIUM.** This component of the spread accounts for investor unfamiliarity with insurance-linked securities. The novelty premium will eventually disappear as investors educate themselves about catastrophe bonds and as transaction structures become more standard. To some degree, this has already happened.
- ❑ **LIQUIDITY PREMIUM.** Catastrophe bonds are relatively illiquid. The illiquidity premium played a very important role when the very first cat bonds were issued. At the time, there was virtually no liquidity in the marketplace, and investors were limited to the buy-and-hold strategy. Over time, however, it has become easier to trade catastrophe bonds. Even immediately before hurricane landfall, when evacuation warnings are issued, it is usually possible to buy and sell securities potentially affected by the hurricane. Initially some structurers of insurance-linked securities have made a special effort to provide liquidity in order to help develop the overall cat bond market. While liquidity is now improving, the bid-ask spreads are still relatively wide and some bonds remain largely illiquid. As the market is growing quickly, both in terms of the number of securities issued and the number of investors, liquidity should continue to improve, reducing the liquidity premium now included in the excess spread.
- ❑ **“SUDDEN-DEATH” PREMIUM.** A cat bond may have the same rating as corporate debt, but there is a very important difference in the timing of default. The default of a corporate bond is usually preceded by the deterioration of the financial condition of the issuer and gradual downgrades by rating agencies. Sudden defaults are rare. Cat bonds, on the other hand, could default with no prior warning or rating agency downgrade. For example, an earthquake could cause an immediate default, resulting in total loss to investors. For some investors, the possibility of a sudden default is unsettling. Certain investors prefer never to see a default in the portfolios, and would sell a security if it is downgraded and chances of default increase. This behaviour is often based on purely psychological factors, with portfolio managers not wanting to be blamed for defaults in their portfolios.
- ❑ **ASYMMETRIC INFORMATION PREMIUM.** This component of the excess spread is present in cat bonds with indemnity-based triggers. Investors in indem-


nity-based cat bonds are at an information disadvantage relative to the insurance company sponsoring the bond. The company has better knowledge than investors of the riskiness of its portfolio of insurance policies.

- **RE-RATING PREMIUM (DISCOUNT).** Sophisticated investors often do not rely on the ratings assigned to cat bonds by rating agencies. Based on their own analysis, investors may choose to not believe ratings for any security and effectively re-rate them by internally assigning their own ratings for the purposes of pricing and risk analysis. This situation is much more common with cat bonds than with other rated securities. Some rating agencies even have caps on ratings assigned to cat bonds. In general, investors tend to believe that cat bonds deserve higher ratings than those assigned by rating agencies. The explanation is that rating agencies, like some investors, might be averse to the situation of sudden default without prior downgrade, and consequently assign ratings to cat bonds based on criteria stricter than those applied to other securities. Another differentiator of cat bonds from other securities is the greater average loss given default (LGD) than for most bonds. Many cat bonds, if defaulted, would likely suffer full default with total loss to investors. Since some investors tend to think that the “real” rating is higher than the one assigned by rating agencies, the excess spread is reduced. In other words, this component of the excess spread, if present, would usually be negative.

It is important to note that, for some catastrophe peril types, the risk during the term of the bond is not uniform. For example, hurricane season in the Caribbean lasts from June till November; the rest of the year, the probability of a hurricane is low. The dependence of risk level on the time period allows us to construct a type of term structure for a catastrophe bond. The non-uniformity of the risk distribution over time has a significant effect on pricing levels in the secondary market.

Because cat bond sponsors usually have the option of reinsuring their risk instead of securitisation, the price levels in the reinsurance market have some effect on the cat bond spreads, in particular the original spreads at issue.

Spreads on cat bonds have been subject to significant volatility. Initially very high, they trended downward until the 2005 hurricane season, when demand level increased. The yields increased in 2005 also because questions were raised about the quality of modelling and analysis provided to investors. The reliability and accuracy of the cat modelling software were questioned, resulting in improvements to the models and reassessment of

Figure 3.9 Satellite image of Hurricane Katrina before landfall

Source: National Oceanic and Atmospheric Administration / US Department of Commerce

Note: Total insured losses from Hurricane Katrina are estimated to be over US\$40 billion, while economic losses are significantly higher.

the catastrophe insurance risk in general. The previously mentioned difficulties encountered by the cat bond market in the second half of 2008 led to the greatest period of volatility and depressed values. This changed in the first half of 2009, when the new collateral structures and the hardening of catastrophe reinsurance markets led to the renewed growth of the market and more stability in pricing.

MARKET STABILITY AND GROWTH

The first loss in a publicly disclosed catastrophe bond was the Kamp Re transaction, in which the risk of a hurricane was transferred to the capital markets investors. Hurricane Katrina in 2005 caused insurance losses of a level that led to the full loss of interest and principal for Kamp Re investors. The loss tested the cat bond market, which prior to Hurricane Katrina had not been known to result in losses to investors. In fact, overall, investors have profited handsomely from catastrophe bonds, with spreads usually being significantly over comparably rated corporate bonds. The default of the cat bonds affected by the bankruptcy of Lehman Brothers as the total return swap counter-party was another difficult test for investors. The market addressed the issues of credit risk by introducing new cat bond structures (see Chapter 7).

The 2004 and 2005 hurricane seasons in the US generated a renewed focus on catastrophe risk management in the insurance and reinsurance industry. The analysis, along with recalibration of catastrophe models, led to the realisation that the risk exposure is far greater than previously believed. This created a strong demand for cat bonds and other capital markets solutions on the part of insurers. The demand was boosted by the limited reinsurance capacity for catastrophe risks.

Hurricane Katrina had an additional impact: the payout of the Kamp Re bond to its insurance sponsor clearly demonstrated that cat bonds could provide reliable protection to insurance companies.

Fixed-income investors are also increasingly interested in catastrophe bonds and other insurance-linked securities. With investors searching for new types of securities to provide diversification and yield, the growing insurance-linked securities marketplace appears more and more attractive.

MORE ON THE SPONSOR AND INVESTOR PERSPECTIVES

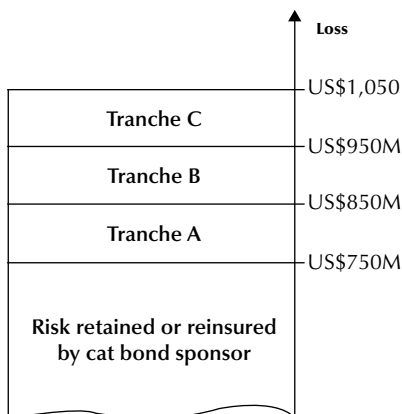
The structure and pricing of a cat bond are an outcome of the process of trying to find a balance between the interests of the sponsor and the investors.

Diversification

A key reason for investors to buy cat bonds is to diversify their investment portfolios. This is true even for the specialised hedge funds that invest only in insurance-linked securities, since other investors obtain diversification by investing in these funds either directly or through the fund-of-funds mechanism. Cat bonds provide investors with a financial instrument weakly correlated with the equity and fixed-income markets, which has led to cat bonds being called zero-beta securities.

The view that there is no correlation between the performance of cat bonds and that of other securities was initially questioned in the aftermath of Hurricane Katrina. While a typical hurricane would not affect financial markets, a very large catastrophic event such as an earthquake in California could have a shock effect on the economy. In these extreme cases, many types of risk suddenly become highly correlated, even if the degree of dependence is very low under normal circumstances. The “zero-beta” view was clearly shown to be invalid by the events of 2008, which uncovered sources of correlation with the markets that had never been appreciated before that period.

While the zero-beta view is incorrect in its application to cat bonds, the

Figure 3.9 Satellite image of Hurricane Katrina before landfall

relatively weak correlation of cat bonds with traditional financial assets is a major source of potential diversification and a strong reason for investors to gain exposure to this asset class. Cat bonds undeniably provide a diversification benefit in addition to affording exposure to a new type of investment.

Within a portfolio of catastrophe bonds and related securities, investors can achieve diversification in a variety of ways. One of them involves building the portfolio with an eye on geographic and peril diversification. Managing portfolios of cat bonds is described in Chapter 16, in the broader context of active management of portfolios comprising various types of catastrophe insurance-linked securities.

Slicing and packaging of risk

A cat bond designed to securitise the risk to an insurance portfolio resulting from a specific natural catastrophe would generally consist of tranches with various degrees of risk. In the example shown in Figure 3.9, if the total loss level exceeds US\$750 million, tranche A is activated. As long as the aggregate loss level remains below US\$850 million, investors in tranche B and tranche C receive interest and principal in full. Since the loss level is above US\$750 million, investors in tranche A suffer the loss of part or entire interest and principal.

To avoid moral hazard, there is usually participation by the sponsor in the excess losses. In structuring terms, this means that not all of the excess risk is reinsured to the SPV, and the sponsor retains a share of potential excess losses.

Figure 3.10 Example of tranches with various types of risk

Tranche 6	– Combo tranche (multi-peril)
Tranche 5	– Hurricane in California
Tranche 4	– Windstorm in the US
Tranche 3	– Hurricane in the US (industry loss trigger)
Tranche 2	– Hurricane in the US (parametric trigger)
Tranche 1	– Earthquake in Japan

Since the tranches have different degrees of risk, they would generally be assigned different ratings, with tranche C as the safest, receiving the highest rating and the lowest spread. It is possible for some tranches to be unrated and others to be rated, in the same transaction.

Another way to slice and package risk is to issue several tranches, with each individual tranche associated with the risk of a specific natural catastrophe in a certain geographic region. Each tranche would have its own trigger; trigger type may even differ from tranche to tranche. A “combo” tranche could also be issued, based on the combination of risks contained in individual tranches. This combination tranche provides diversification to investors unwilling or unable to achieve it on their own. Figure 3.10 provides an example of such a structure.

The Successor cat bond issued by Swiss Re in 2006 is a good example of this structure. The Successor programme placed US\$950 million of principal-at-risk variable-rate notes, transferring to investors the risks of North Atlantic hurricane, European windstorm, California earthquake and Japanese earthquake in individual and multi-peril tranches.

Another pioneering transaction brought to the market by ABN Amro in 2006 was structured as a collateralised debt obligation (CDO) from the very beginning. In fact, it was the first publicly rated CDO of natural catastrophe risk. The CDO offered to investors was based on the cat bonds with industry loss triggers sponsored by the Catlin Group. The least risky tranche of the CDO was then rated AA by Standard & Poor's. Higher ratings open up a new universe of investors who otherwise would have no interest in catastrophe insurance-linked securities. The negative connotation of the term

CDO has led to renaming this type of security collateralised risk obligation (CRO). A managed CRO structure was introduced by Nephila Capital in the Gamut transaction developed by Goldman Sachs in 2007. At this point, it is unclear whether CRO structures will be actively used in the future.

Types of sponsor


Catastrophe bonds are generally sponsored by insurance or reinsurance companies. However, corporations can also get protection against natural catastrophe losses by going directly to the capital markets. Tokyo Disneyland's securitisation of earthquake risk in Japan provides an example of a non-insurance company bypassing the insurance marketplace and going directly to the capital markets to obtain cat protection.

Many believed that in the future cat bonds would be issued only on behalf of reinsurance companies. This view was based on its being seemingly more efficient for primary insurance companies to reinsure their risk as opposed to sponsoring cat bonds. Reinsurance companies would then accumulate all the risk, and transfer a part of it to the capital markets. This has not happened and we do see cat bonds issued directly by insurance companies. One of the reasons is the credit risk involved in catastrophe reinsurance. Reinsurance companies are particularly exposed to the risk of natural catastrophes, and might default on their obligations should such an event happen. Cat bonds, on the other hand, provide a fully collateralised protection with little exposure to credit risk.

From the point of view of an investor, the identity of the sponsor of a non-indemnity cat bond is largely irrelevant, with the analysis focused on natural catastrophe modelling performed by the same cat modelling firms as would be modelling insurance company books of business.

Investor types

While the first major investors in cat bonds were reinsurance companies, now they represent only a small percentage of the overall investor base. A number of specialised hedge funds have been formed for the sole purpose of investing in insurance-linked securities. These funds often possess superior expertise and drive the pricing of cat bonds both at issue and in the secondary markets. In addition, many other investors such as pension funds have invested in cat bonds. The number of investors in insurance-linked securities and the total capital committed to this asset class continue to grow.

Figure 3.11 Catastrophe modelling technology

MODELLING PROPERTY CATASTROPHE INSURANCE RISK

The reason for including risk analysis in cat bond offering documents is that investors do not have the means to assess default probabilities on their own. This is the case in part because most of them do not possess expertise in determining the likelihood of natural catastrophes and resultant insurance losses. The other reason, applicable to indemnity-type transactions, is that detailed information on the exposure by geographic location is not provided, making it impossible for investors to determine exact default probabilities even if they had superior expertise in analysing insurance risk of hurricanes and earthquakes.

Specialised catastrophe modelling firms play a critical role in the securitisation of insurance catastrophe risk. The modelling software provides the only objective way to analyse the probability of default. It is also the only way for rating agencies to assess the risk and be able to assign a rating to these securities.

The modelling generally includes two components. First, a probabilistic analysis of specific types of natural catastrophes is performed for a certain geographic area. For example, the model could simulate hurricanes in Florida or typhoons affecting Japan. The second step involves assessing the financial impact that these natural catastrophes would have on the portfolio of insurance policies held by the sponsoring insurance company. This assessment is also probabilistic. The final output of the model is the probability distribution of the insured losses, which takes into account not only policy conditions and limits, but also the reinsurance structure in place.

The damage module is based on structuring engineering input. Its function is to take a specific catastrophe scenario and superimpose it onto a portfolio of insurance policies being analysed. The damage calculator takes individual exposures such as insured properties and probabilistically estimates the damage caused by the catastrophe under the scenario, taking into

account such parameters as policy limits and deductibles. The output is the insured loss that the company would have to pay out under the scenario. The model runs a large number of scenarios and generates a set of catastrophic losses and probabilities associated with them. Chapter 4 contains more detailed treatment of modelling catastrophe risk.

There are only three major recognised independent providers of modelling services for insurance catastrophe exposure. The three companies, AIR Worldwide, EQECAT and Risk Management Solutions (RMS), are primarily software developers for the property-casualty insurance industry. While the RMS model is the most widely used in the industry, AIR is currently leading in providing consulting analytical services for structuring cat bond transactions.

The output of catastrophe modelling software includes the data necessary to construct an exceedance probability curve. The exceedance probability curve could be used for structuring and pricing a cat bond. In structuring, it would help determine the trigger level to provide the needed protection to the insurance company. In pricing, the exceedance probability curve is used to provide a probabilistic look at exceeding the trigger level (that is, bond default) that determines the bond price.

TRENDS AND EXPECTATIONS

The catastrophe bond marketplace is growing and will continue to do so, along with other capital markets mechanisms for transferring catastrophe insurance risk. We are witnessing both an increase in cat bond issuance and growth in the total capital committed to this asset class. Some of the reasons for the growth and its drivers are as follows.

- ❑ The insurance-linked securities market has finally reached the critical mass needed to make cat bonds a solution always to be considered in evaluating available options in the transfer of insurance catastrophe risk.
- ❑ The 2004 and 2005 hurricane seasons have led to an increased emphasis on catastrophe risk management. This emphasis has been both internal and external, stimulated by increased scrutiny by the rating agencies and regulators. It has resulted in a demand for additional catastrophe-risk-bearing capacity that is not met by traditional reinsurance mechanisms.
- ❑ The recalibration of catastrophe models post-Katrina has led to the realisation that the insurance industry is exposed to much greater risk of natural catastrophes than previously thought.

- ❑ The second half of 2008 was the greatest test of the viability and future prospects of the market. The bankruptcy of Lehman Brothers led to the default of cat bonds for which Lehman served as the total-return-swap counterparty. Besides the counterparty risk, these events revealed structural weaknesses in the way collateral arrangements had been made in the standard cat bond structures. Ultimately, however, the market has emerged from this debacle stronger, as the weaknesses were addressed in new structures and all other potential weak points carefully examined.
- ❑ The depressed values of cat bonds in 2008 caused by the forced selling of cat bonds by multi-strategy hedge funds made the low-correlation (low-beta) argument slightly weaker, to some degree reducing the diversification value of cat bonds. However, it also highlighted the advantages of this asset class: the multi-strategy funds faced with redemptions were selling cat bonds because they held value better than the great majority of other asset classes.
- ❑ The educational process in the insurance industry has led to better understanding of cat bonds and other risk-linked securities, allowing insurance and reinsurance companies to see the advantages of the securitisation approach.
- ❑ Investors, too, have become better educated about catastrophe bonds and the benefits of diversification provided by these securities. The number of investors in risk-linked securities is growing, including the hedge funds focused exclusively on insurance risk.
- ❑ Structuring of catastrophe bonds has become more standardised, making the process easier for the sponsors and the analysis more straightforward for investors.
- ❑ The cost of issuance of catastrophe bonds has gone down, due to the standardisation of cat bond structures, the use of multi-year bond terms to spread the cost of issuance over a longer period of time, and shelf registration.
- ❑ With the growth in the number of cat bonds issued and in the total investor capital, the secondary market for cat bonds and similar securities is growing, too, resulting in greater liquidity. This, in turn, creates greater opportunities for active management of investment portfolios including cat bonds.

Other important developments that will affect the future of the market are the following.

- ❑ Innovation is continuing, resulting in new products or modifications of the old products to better suit the needs of both issuers (sponsors) and investors.
- ❑ There has been some movement away from indemnity-based towards parametric index triggers, with bond default not depending on the actual losses of a specific insurance company. Many investors are no longer willing to be at an informational disadvantage and demand that default triggers and payout be based on a more objective index.
- ❑ With the movement away from indemnity-based triggers, basis risk is becoming a growing concern for the sponsors of catastrophe bonds. The risk that cat bonds would turn out to be an ineffective hedge and will not provide protection when expected is necessitating better modelling and trigger choices.
- ❑ The development of new parametric triggers that have the ability to further reduce basis risk of the sponsors is an ongoing process and will likely lead to the greater use of these new triggers at the expense of the indemnity and standard industry loss triggers. The ability to address the issue of basis risk can expand the universe of sponsors and lead to market growth.

Securitisation of new types of insurance risk, including liability insurance, will probably grow and has a potential to become a viable alternative to reinsurance for some extreme catastrophic events. It is also expected that insurance securitisation will move beyond very low-frequency/extreme-severity events and will involve higher-frequency insurance risk.

¹ Transformed by Swiss Re on behalf of Zurich American Insurance.

Modelling Catastrophe Risk

THE CHALLENGE OF MODELLING CATASTROPHE EVENTS

The very last painting by Salvador Dali was titled *The Swallow's Tail – Series on Catastrophes*. Dali was greatly interested in the catastrophe theory developed by the French mathematician René Thom, and referred to it as “the most beautiful aesthetic theory in the world”. Thom’s catastrophe theory describes how small changes in parameters of a stable nonlinear system can lead to a loss of equilibrium and dramatic, on the level of catastrophic, change in the state of the system. Thom described equilibrium topological surfaces and corresponding discontinuities that exist under certain conditions. An equilibrium state is associated with the minimum of its potential function; according to the catastrophe theory, a phase transition or a discontinuity can be associated with only a limited number of stable geometric structures categorising degenerate critical points of the potential function. *The Swallow's Tail* includes two of the so-called elementary catastrophes taken directly from Thom’s graphs: the swallowtail and cusp geometries. Dali was captivated by the catastrophe theory, especially after he met Thom. *Topological Abduction of Europe – Homage to René Thom*, an earlier painting by Dali, even reproduces in its bottom left corner the formula describing the swallowtail elementary catastrophe geometry.

There have been numerous attempts to apply the catastrophe theory to describing and predicting physical events. Returning from art to science, we are faced with the challenge of assessing the frequency and severity of natural and manmade catastrophes that can lead to massive insurance losses. The challenge is daunting, and developing a model to accomplish this goal is a very practical task – with no surrealistic elements, even if the results of catastrophes can often appear surreal. This chapter introduces important concepts in modelling catastrophic events for the purpose of analysing insurance risk securitisation. Issues examined here provide an understanding of why modelling catastrophe risk is essential and why it is often so challenging.

Predicting the unpredictable

Catastrophic events are impossible to predict. The only way to analyse these events and their impact on insured losses is within a probabilistic framework. Catastrophe modelling has evolved in recent decades: its role in quantifying insurance risk is critical and credible. The credibility of the modelling tools continues to grow as they incorporate more and more of the latest scientific research on catastrophic events and the insurance-specific data that determines the impact of the catastrophes on insurance losses.

IMPORTANCE OF CATASTROPHE MODELLING TO INVESTORS

Wherever the payout on insurance-linked securities is tied to the possible occurrence of insured catastrophe losses, catastrophe modelling is the most important tool for investors in analysing the risk of the securities and determining the price at which they would be willing to assume this risk.

Superior ability to model insurance risk of catastrophic events is a source of competitive advantage to investors in securities linked to such risk. This ability can serve as an important differentiator and an indispensable tool in a market that remains inefficient and suffers from the problem of asymmetric information and general information deficiency.

The chapter on catastrophe bonds provided a brief overview of the structure of the models used in analysing the insurance risk of property catastrophe securitisations; it also examined important outputs such as exceedance curves that specify probabilities of exceeding various loss levels. It is equally important to understand inputs to the models.

The seemingly straightforward task of understanding the results, such as interpreting the risk analysis included in the offering documents for cat bonds, is actually the most important and the most challenging. If the modelling software is a complete black box to an investor, any analysis of its output is limited and deficient. Not understanding the modelling tools also detracts from the usefulness of the sensitivity analysis that might be included in the offering documents; it makes it difficult to make any adjustments to improve on what is included in the documents.

It is unrealistic for most investors to become familiar with the inner workings of catastrophe modelling software to get a better insight into the risk involved in insurance-linked securities. The cost-benefit analysis does not justify developing such expertise in house. Only true specialists can afford this luxury. However, it is beneficial to any investor in catastrophe insurance-linked securities to be familiar with the basic methodology of modelling catastrophe risk. This, at the very least, will allow

investors to interpret the data in the offering circulars on a more sophisticated level.

MODELLING CATASTROPHE INSURANCE RISK OF INSURANCE-LINKED SECURITIES

The chapter on catastrophe bonds provided an overview of the modern catastrophe modelling technology and described the main modules of a catastrophe modelling software provided by the three recognised independent providers of insurance catastrophe modelling services, AIR Worldwide, EQECAT and Risk Management Solutions (RMS). The chapter also introduced concepts such as exceedance probability curve and return period, and included a summary of sensitivity analysis and stress testing that can be performed in evaluating insurance-linked securities.

The output of a catastrophe model is based on thousands or even millions of years of simulated natural events and their financial impact on a given insurance portfolio. This output can then be used to determine the probability distribution of cashflows for a catastrophe bond or another security linked to the risk of catastrophic events.

In fact, the modern models are not limited to natural catastrophes: models of manmade catastrophes have also been developed. For example, terrorism models have been developed to model the risk of catastrophe losses resulting from such acts.

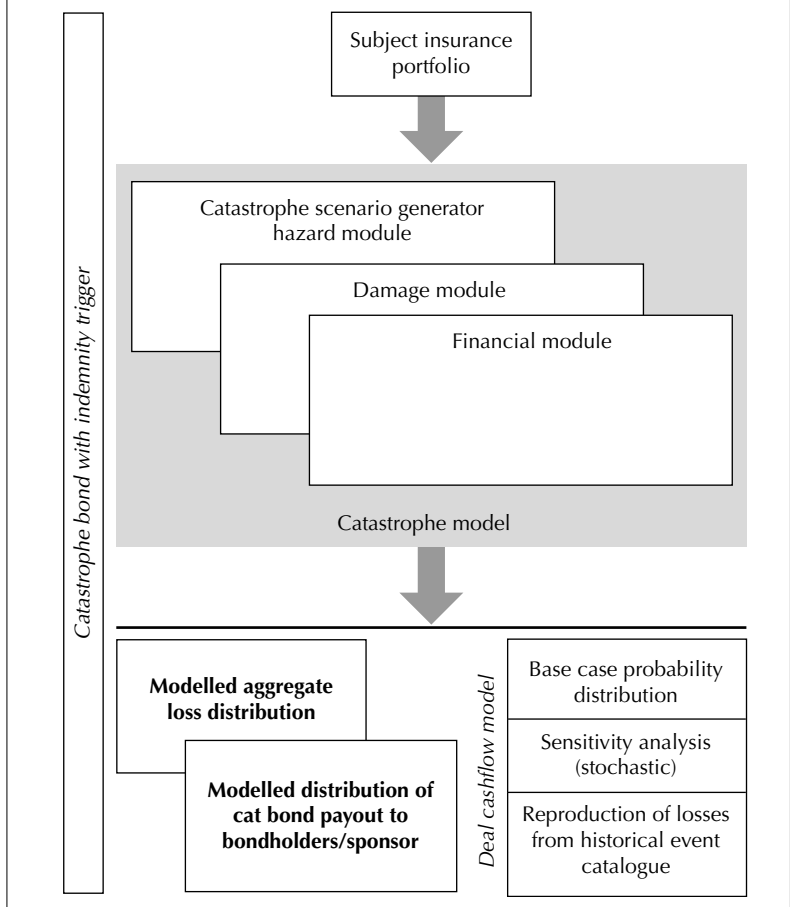
In this chapter, more information on the practical ways to model the cat risk of ILS is added, along with a description of the available modelling tools, their benefits and their limitations. First, however, the basics of the science of natural catastrophes are described, since they form the framework for the generation of catastrophe scenarios used by these software tools.

THE SCIENCE OF CATASTROPHES

It is neither possible nor necessary for an investor to have in-house experts on the actual science underlying catastrophe models; basic understanding, however, at the very least allows us to ask the right questions and to bring a degree of transparency to the black-box view of the models.

Seismology is the study of earthquakes and the physical processes that lead to and result from them. In the broader sense, it is the study of earth movement and the earth itself through the analysis of seismic waves. Earthquake prediction *per se* is not possible, but it is possible to identify probabilities of earthquakes of specific magnitude by geographic region; in

Figure 4.1 The use of insurance catastrophe modelling software for creating a probabilistic deal cashflow model for a catastrophe bond with an indemnity trigger



some cases, there are precursors that might be useful in short-term forecasting as well.

Climatology and meteorology are the study of weather and atmospheric conditions, with the latter focused on the short-term analysis of weather systems and the former on the long-term analysis of weather patterns and atmospheric phenomena. The study of catastrophic weather events such as hurricanes is a specialised branch of this science. In recent years, significant progress has been made in understanding the dynamics of weather-related catastrophes, and in assessing both long-term and short-term probabilities of such events.

Structural engineering and several related fields permit the analysis of damage to physical structures given the occurrence of a specific natural catastrophe. This analysis is important for assessing insurance losses that can result from a catastrophe such as hurricane or earthquake.

Epidemiology and medicine offer yet another example of study of catastrophes, examining pandemic-type catastrophe events and their impact on the population.

Manmade catastrophes are as difficult to predict as those caused by nature; disciplines ranging from structural engineering to political science can provide input into creating a probabilistic model of this type of catastrophic events.

EARTHQUAKE FREQUENCY AND SEVERITY

A simple relationship between earthquake frequency and magnitude is described by the Gutenberg–Richter law. It states that, for a given long period of time in a certain region, the number N of earthquakes of magnitude M or greater follows the power law

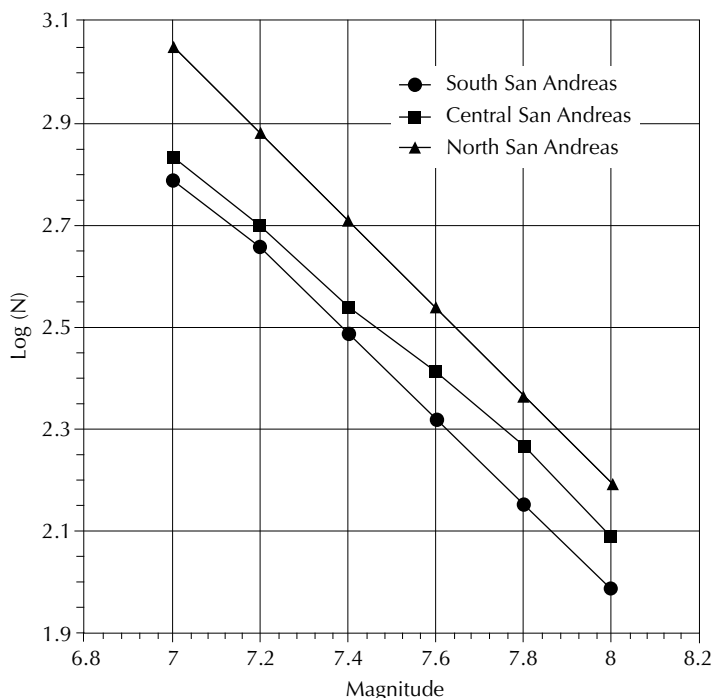
$$N(M) = 10^{a-bM},$$

which can alternatively be written as $\log N(M) = a - bM$, where a and b are constant. b usually, but not always, falls in the range between 0.8 and 1.2. This relationship, specifying that an earthquake magnitude has a left-truncated exponential distribution, holds surprisingly well for many territories and earthquake magnitudes. It can be used to obtain rough estimates of the probability of earthquakes, even of magnitudes not observed, based on the observations of earthquakes of other levels of magnitude.

Another important relationship is the Omori–Utsu law,¹ which describes the aftershock frequency of an earthquake. According to the Omori–Utsu law, the rate of aftershocks decays after the main shock as

$$n(t) = \frac{K}{(t+c)^p},$$

where $n(t)$ is the aftershock frequency at time t after the main shock, and K , c , and p are constant. The c constant is the time-offset parameter describing the deviation from the power law immediately after the main shock. The Gutenberg–Richter law can be used to describe the distribution of aftershocks by magnitude, which shows that the aftershock magnitude decay can also be described by a power law. The Reasenber–Jones model combines the Guttenberg–Richter and Omori–Utsu laws to describe the intensity both of the main shock of an earthquake and its aftershocks.

Figure 4.2 Gutenberg–Richter law: San Andreas fault

The data is presented over a hypothetical 10,000 year period for three sections of the San Andreas Fault

Source: US Geological Survey (T. Parsons)

According to Bath's Law, in an earthquake, the difference in magnitude between the main shock and its strongest aftershock is constant and independent of the earthquake magnitude. All of these models should be considered in a probabilistic framework.

It is important to note that the scientific definition of aftershocks, according to which they can happen years or decades after the main shock, differs from the insurance definition, which has a very narrow time range for what constitutes an earthquake event. Insurance-linked securities such as catastrophe bonds follow the same narrow definition of an earthquake, with aftershocks having to fall within a defined short period of time after the main shock; otherwise, an aftershock might be considered a separate earthquake event, and in that case it might have different coverage terms, it might not be covered at all, or it might trigger second-event coverage.

The basic phenomenological laws such as the Gutenberg–Richter and

Table 4.1 Examples of earthquake scales

	Richter	Modified Mercalli
Description	Logarithmic scale to measure the amount of seismic energy released by an earthquake	Based on subjective description of damage and feeling of shaking; value changes with distance from hypocentre
Range and effects	From 2.0 (recorded by instruments not felt) to 9.9 (great devastation in areas up to several thousand miles across if epicentre close to surface). 10.0 and greater never observed.	From I (neither felt nor caused noticeable damage but recorded by instruments) to XII (catastrophic damage with almost everything destroyed)

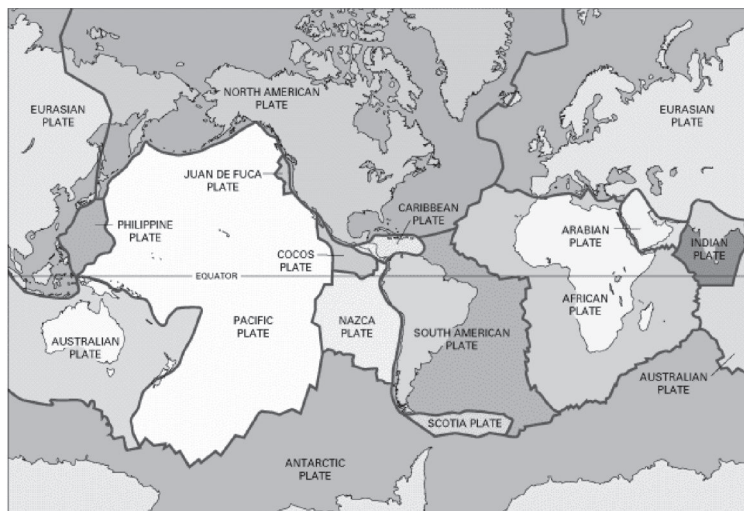
Omori–Utsu relationships are more accurate than their simple form would suggest. However, such simple laws are obviously insufficient for modelling earthquakes, and several more sophisticated models have evolved for this purpose.

EARTHQUAKE LOCATION

The vast majority of earthquakes occur on tectonic plate boundaries; though some, typically smaller ones, do occur within the plates. Earthquakes within the tectonic plates usually happen in the zones of fault or weakness, and occur only in response to pressure on the plate originating from its boundary. The three categories of tectonic plate boundaries are spreading zones, transform faults and subduction zones, each of which can generate its own type of earthquake. Most spreading zones and subduction zones are in the ocean, while transform faults can occur anywhere and are among the best studied.

A global map of tectonic plates is presented in Figure 4.3, overleaf; it shows the main tectonic plates and the boundary lines between them.


The hypocentre, where a rupture happens, is typically not very deep under the earth’s surface for transform faults. In other words, the distance between the hypocentre and the epicentre is relatively small. Compressional and dilatational movements tend to follow straight patterns, at least for “simple” earthquakes such as those that involve limited changes to the original earthquake slip. The study of faults plays a major part in determining the probability distributions of earthquakes in different areas.

Figure 4.3 Map of tectonic plates

Source: US Geological Survey

Seismic hazard maps illustrate the distribution of earthquake shaking levels that have a certain probability of occurring. Figure 4.4, opposite, shows the US national seismic hazard map that displays shaking levels, expressed as peak ground acceleration (PGA), at the probability level of 2% over the period of 50 years. Other maps developed by the US Geological Survey (USGS) correspond to the 5% and 10% probability of exceedance over the 50-year period. The map shown was developed in 2008; the USGS produces a fully revised version of the national seismic hazard maps approximately once every six years. The national seismic hazard maps are important in insurance catastrophe modelling even if the modellers disagree with the methodology used in developing the maps: the maps form the basis for many building codes, which in turn determine the level of property damage in case of an earthquake of a certain magnitude.

The two main types of earthquake models are fault- and seismicity-based. The fault-based models rely on fault mapping; each known fault or fault segment has a statistical function associated with the recurrence time for earthquakes of specific magnitude. In the simplest case, it is assumed that following an earthquake at a fault, stress on the fault has to be “renewed” by the tectonic processes until the next earthquake occurs. This view, while fully stochastic, implies a certain degree of regularity of earthquakes that

Figure 4.4 US national seismic hazard map

The shadings represent the levels of horizontal ground shaking (peak horizontal ground acceleration) that have a 2% exceedance probability over a 50-year period. Shaking is expressed as a percentage of g (the acceleration due to gravity).

Source: US Geological Survey

leads to quasi-periodicity of earthquake occurrence. This is why fault-based models are also referred to as renewal models. Poisson, Weibull, gamma or lognormal distributions can be used in modelling time between earthquakes, even though other arrival process distributions are sometimes utilised as well. The Poisson renewal process, with an exponential distribution of recurrence times, is the simplest but probably least accurate. In its simplest form the Poisson fault-based model is time-independent. In contrast to the fault-based models, seismicity-based models assume that observed seismicity is the main indicator of the probability of future earthquakes. The use of the Gutenberg–Richter law or a similar relationship then allows the observed frequency of small earthquakes to be used for estimating earthquakes of greater magnitude. This approach does not require information on the faults or even knowledge of their existence; it overcomes a drawback of fault-based models, which can fail because many faults are not yet mapped correctly, and some are not mapped at all. Seismicity-based models are also called cluster models: the occurrence of several smaller earthquakes might signify the coming of a bigger one. Renewal processes can be used also for describing clustering events. Aftershock models allow us to project past seismicity forward to arrive at a time-dependent proba-

bility distribution of earthquakes at a specific location. The fault- and seismicity-based models are not mutually exclusive: elements of both are employed in modelling, in particular for the better-researched faults for which there is also more extensive seismicity data available.

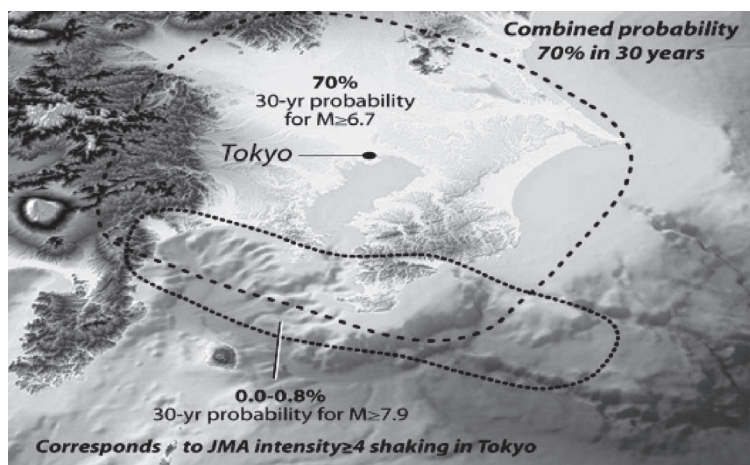
Some parts of the world have high levels of earthquake-related insurance risk. They combine greater probability of earthquakes, due to being situated on or close to a fault line, and the concentration of insured risk exposure. All of Japan and part of California are examples of such high-risk areas.

Japan is located in a very seismically active area and has very high density of population and insured property. Earthquakes in Japan have claimed many lives and caused significant property damage. The growth in population and property has led to the situation whereby a repeat of one of the historically recorded earthquakes would now result in enormous losses. Estimates of the overall (not only the insured) cost of a repeat of the great 1855 Ansei-Edo earthquake today go as high as US\$1.5 trillion. Tokyo sits at the junction of three tectonic plates: it is located on the Eurasian plate; while not far from the city the Pacific tectonic plate “subducts” from the east, and the Philippine Sea tectonic plate “subducts” from the south. Of particular concern is the plane fragment under the Kanto basin, detached from either the Pacific or the Philippine Sea tectonic plate, whose position could lead to a large-magnitude earthquake in the already seismically active region.

Japanese earthquakes have been modelled very extensively, but there remains a significant level of uncertainty as to the probability distribution of their frequency and severity. This particularly high level of uncertainty has to be taken into account in any analysis of earthquake risk in Japan.

It has been said that the occurrence of a large-magnitude earthquake in a densely populated area in California is a question of not if but when. The San Andreas Fault is situated where the North American tectonic plate and the Pacific tectonic plate meet, with the North American plate moving southward and the Pacific plate northward. The fault, shown on Figure 4.6 on page 66, goes almost straight through San Francisco, with the city being on the North American plate, slightly to the east of the San Andreas Fault. Los Angeles is also situated dangerously close to the fault line, but is located to the west of it on the Pacific tectonic plate. San Andreas is a transform fault; transform faults tend to produce shallow earthquakes with the focus close to the surface.

A number of studies have concluded that there is a high probability of a major earthquake at the San Andreas fault system, in particular in its southern part, where stress levels appear to be growing and where there has

Figure 4.5 Probability of high intensity earthquake affecting Tokyo

Source: US Geological Survey (based on Japanese government data)

not been a major earthquake in at least three centuries. The conclusion that the southern part of the fault has a higher probability of a major earthquake is not universally accepted. There is an agreement that all areas along the fault, including San Francisco, which experienced a major earthquake in 1906, are at significant risk.

MORE ON EARTHQUAKE MODELLING

A numerical simulation approach has been used for modelling earthquake parameters. The nature of the earthquake phenomenon and its inherent uncertainty invites the probabilistic approach, and simulation is the natural way to implement it. Models have been developed for describing ground motion, stresses at the faults, fault dimensions, rupture velocities and many other parameters.

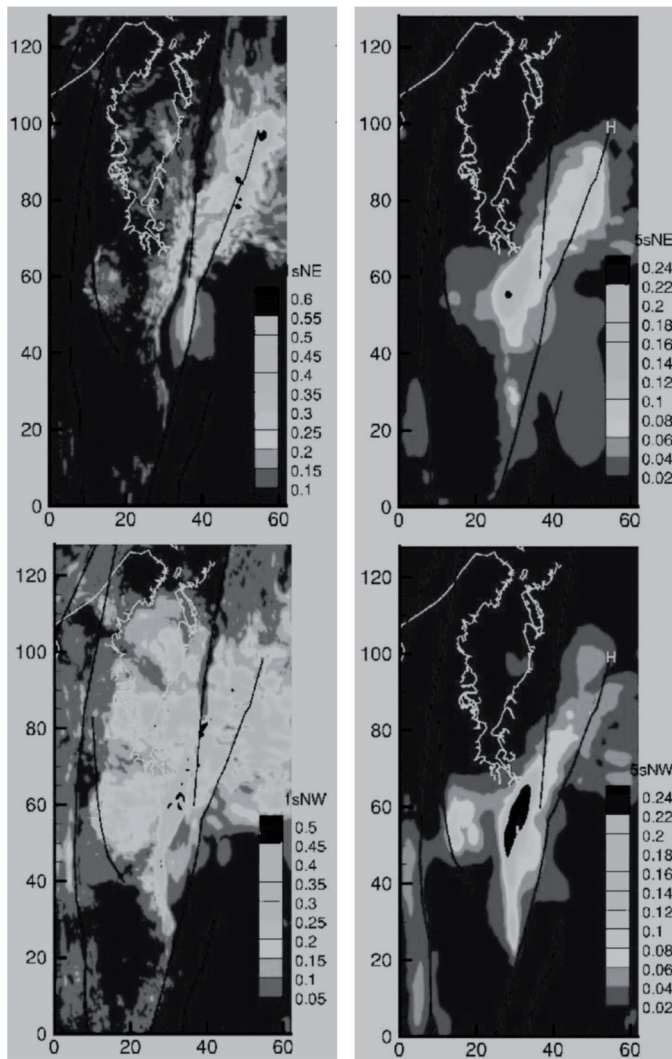
The sheer number of unknowns and random variables involved in simulating earthquakes leads to attempts to simplify the problem by focusing on only major factors affecting the development of earthquakes, and by using phenomenological laws in place of direct simulation for some variables. The results have been mixed. While every one of the existing models and approaches is incomplete, relies on many simplifying assumptions and could be easily criticised, there has not yet emerged a way to adequately simulate such complex natural phenomena as tectonic developments and earthquakes.

Figure 4.6 San Andreas Fault

Source: US Geological Survey

Even though the numerical simulation approach is generally the best to portray the behaviour of complex systems, incorrectly specifying some of the variables or the interdependences among the variables can lead to incorrect results. Even simpler approaches, by necessity neglecting interdependence of some of the variables involved, are very challenging to implement. Fitting distributions to variables such as the recurrence times of major earthquakes is a common approach. It still leaves a lot of room for uncertainty even as far as the choice of the probability distribution to be

Figure 4.7 Simulating earthquakes: ground motion in Santa Clara Valley, California, and vicinity from M6.7-scenario earthquakes and greater



Pseudo-spectral acceleration (PSA) (in units of g, 5% damped) for a M6.8-scenario earthquake on the Calaveras CN fault segment with epicentre near Danville (H). Left column 1-sec period, right column 5-sec period. Top row northeast component, bottom row northwest component.

Source: Earthquake Hazards Program, USGS (Harmsen *et al.*)

fitted. As an example, Weibull distribution can be used to simulate earthquake occurrence times in the following way

$$P(t, t_0) = 1 - \exp \left[\left(\frac{t_0}{\tau} \right)^\beta - \left(\frac{t}{\tau} \right)^\beta \right]$$

expressing the cumulative probability of an earthquake happening at time t after the last earthquake, conditioned on there having not been an earthquake for a period of time t_0 since the last earthquake.² Parameters τ and β are fitted to the distribution based on available data.

Epidemic-type aftershock sequence (ETAS) models are the most common of the aftershock models mentioned above. They assume that each daughter earthquake resulting from a parent earthquake has its time of occurrence and magnitude distributed randomly but generally based on the Gutenberg–Richter and Omori–Utsu laws. Each daughter earthquake is a parent to the next generation of earthquakes. If the first-generation aftershock is greater in magnitude than the main shock, it becomes the main shock, and the shock previously considered to be the main shock becomes a foreshock. The branching aftershock sequence (BASS) model further imposes Bath's Law in a modified form for the generation of earthquake sequences. Simulations based on the BASS model are often unstable; this practical difficulty can be overcome by imposing additional constraints on simulations. BASS models are seen as providing a better description of aftershock sequences than the standard ETAS models.

A superior approach (though harder to implement) is not to impose a specific probability distribution on the recurrence time variable, but instead to simulate the physics of fault interaction, reflecting the correct topology and process dynamics. The earthquake recurrence times are then the output of that simulation process and do not follow any formulaic distribution.

The models are evolving, and the ultimate goal is to create a complete model of earthquake generation based on the simulation approach. Advances in geophysics and computing make it possible to move closer to this goal. Creating a complete earthquake generation model requires simultaneous simulation of many interrelated processes involved in earthquake generation.³ Large-scale supercomputer simulations are opening doors to creating models that incorporate the latest advances in earthquake physics and physical observations related to specific faults. Results of research coming from the Earth Simulator supercomputing project and other institutions have already been sufficiently valuable to be reflected in some

modelling software used to analyse the risk embedded in insurance-linked securities.


TSUNAMIS

Tsunamis are caused by underwater seismic events such as regular earthquakes, volcanic explosions and landslides. They can also be caused by meteorites or underwater nuclear explosions. Since the causes of tsunamis are usually earthquakes, the study of tsunamis is closely related to earthquake science. Mapping potential earthquake locations and estimating probability of earthquakes of various magnitude at these locations is an important part of the tsunami threat analysis. Another part is estimating the impact of a tsunami caused by an earthquake with known location, magnitude and other characteristics.

Tsunami modelling involves three parts corresponding to the three stages of a tsunami: wave generation, propagation and inundation. Propagation modelling attempts to produce stochastic scenarios of tsunami waves' speed, length, height and directionality. (Even though tsunami waves spread in all directions, there is often one direction that exhibits tsunami beaming, or the higher wave heights.)

Modelling of run-up, which is a term used to describe the level of increase in sea level when the tsunami wave reaches shore, requires good knowledge of underwater topography close to shore. Far-field tsunami wave trains might result in greater inundation than waves of the same run-up heights

Figure 4.8 Map of major recorded tsunami events (epicentres)



Source: National Geophysical Data Center, National Oceanic and Atmospheric Administration (NOAA)

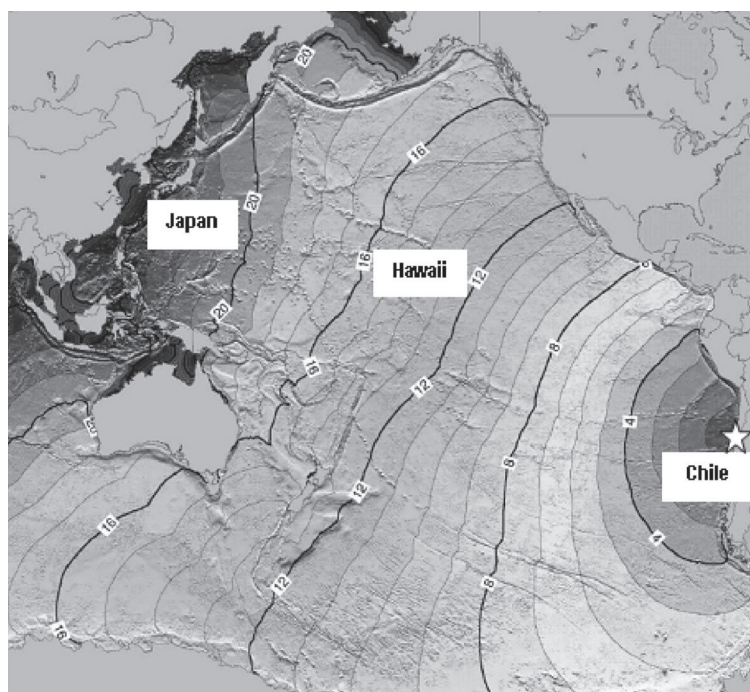
generated by an underwater earthquake or landslide located close to the area of inundation.

A number of models for simulating tsunami events have been developed, and to a significant degree validated. Databases of pre-computed scenarios have been created for such tsunami-prone areas as Hawaii and Japan. High-resolution models are extremely useful in estimating an impact of a tsunami on insured properties.

HURRICANES

Hurricanes represent the main natural catastrophe risk embedded in insurance-linked securities such as catastrophe bonds. A broader term, cyclone, includes both tropical cyclones (hurricanes, typhoons, tropical storms and tropical depressions) and extratropical cyclones, such as European wind-

Figure 4.9 Simulated travel map for a hypothetical magnitude 9.2 underwater earthquake off the coast of Chile with wave propagation across all of the Pacific Ocean



Source: National Geophysical Data Center, NOAA (in collaboration with ICG/PTWS)

storms and Northeasters. North Atlantic hurricanes are the main cyclone risk transferred to investors in insurance-linked securities, followed by European windstorms.

The terminology is not consistent even within the same geographical region. Table 4.2, overleaf, displays the classification based on the criteria established by the US National Oceanic and Atmospheric Administration (NOAA). Hurricanes in the Northwest Pacific are usually called typhoons, while in the southern hemisphere all tropical storms and hurricanes are referred to as cyclones.

A number of cyclone scales are in existence to classify cyclones by their strength. Wind speed is the most important parameter used in the classification systems, but other parameters are used as well. The scales vary by the way they measure storm strength and by which oceanic basin is being considered.

The hurricane risk in insurance-linked securities is most often that of hurricanes striking the US, in particular the hurricanes originating in the Atlantic Ocean. Hence the description below is US-centric; and for this reason the terminology and analytical tools described here are primarily those developed by NOAA and in particular its National Hurricane Center. While the terminology and some of the characteristics of the hurricanes differ around the world, the example of the North Atlantic hurricanes provides a good general illustration, and most of its elements can be applied to cyclones in other parts of the world. In addition, North Atlantic hurricanes are arguably the best researched and documented, with numerous models having been developed for their analysis.

Some of the scales used around the world include the Beaufort wind scale (initially developed for non-hurricane wind speeds but now extended to include five hurricane categories), Dvorak current intensity (based on satellite imagery to measure system intensity), the Fujita scale or F-scale (initially developed for tornadoes but now also used for cyclones), the Australian tropical cyclone intensity scale (similar to the expanded part of the Beaufort scale) and the Saffir–Simpson hurricane scale. The last of these is the primary scale used by NOAA; it divides hurricanes into five distinct categories outlined in Table 4.3 on page 73. In the description of the effects of a hurricane, this scale uses the damage characteristics most appropriate for the US. When applied to categorising hurricanes in other parts of the world, only the level of sustained wind speeds would normally be used.

One of the criticisms of the Saffir–Simpson Hurricane Scale has been the inclusion of specific references to storm-surge ranges and flooding refer-

Table 4.2 Cyclone classification (current NOAA definitions)

Cyclone type or stage of development	Criteria
Tropical depression (development)	The formative stages of a tropical cyclone in which the maximum sustained (1-min mean) surface wind is less than 34 kt (39 mph or 18 m/s)
Tropical storm	A tropical cyclone in which the maximum sustained surface wind (1-min mean) is 34 kt (39 mph; 18 m/s) or greater but less than 64 kt (74 mph; 33 m/s)
Hurricane	A tropical cyclone in which the maximum sustained surface wind (1-min mean) is at least 64 kt (74 mph or 33 m/s)
Major hurricane	A hurricane classified as Category 3 or higher, with maximum sustained surface wind (1-min mean) of at least 96 kt (111 mph or 50 m/s)
Tropical depression (dissipation)	The decaying stages of a tropical cyclone in which the maximum sustained surface wind (1-min mean) has dropped below 34 kt (39 mph or 18 m/s)
Extratropical cyclone	A tropical cyclone that has been modified by interaction with a non-tropical environment, and whose primary energy is baroclinic. There are no wind-speed criteria, and maximum winds may exceed hurricane force.
Subtropical depression	A low-pressure system that develops over subtropical waters and initially may have a non-tropical circulation, but some elements of tropical cyclone cloud structure are present. Surface winds are below 34 kt (39 mph or 18 m/s)
Subtropical storm	Same definition as subtropical depression except that the wind is at least 34 kt (39 mph or 18 m/s). Maximum winds may exceed hurricane force.

Source: NOAA

ences. Parameters such as the topographic profile of the coastline where a landfall happens, forward speed and size of the hurricane at landfall all affect storm-surge levels and can put them outside the range expected based purely on wind speeds. Hurricane Ike in 2008 is an example of such inconsistency. To address this criticism, in 2009 NOAA implemented the Saffir–Simpson Hurricane Wind Scale (the word Wind is added to distin-

guish the two scales), which does not have specific references to the level of storm surge and includes an updated description of the damage effects. While currently considered experimental, it is likely that the new scale will become the main hurricane classification tool in the US. Table 4.4 provides the description of the categories in the 2009 Saffir–Simpson Hurricane Wind

Table 4.3 NOAA Saffir–Simpson Hurricane Scale (based on original Saffir–Simpson scale with minor modifications)

Hurricane category	Sustained wind speed	Effects
1	74–95 mph (64–82 kt or 119–153 km/hr)	Damage primarily to shrubbery, trees, foliage and unanchored mobile homes. No real damage to other structures. Some damage to poorly constructed signs. AND/OR: storm surge 4 to 5 feet above normal. Low-lying coastal roads inundated, minor pier damage, some small craft in exposed anchorage torn from moorings.
2	96–110 mph (83–95 kt or 154–177 km/hr)	Considerable damage to shrubbery and tree foliage; some trees blown down. Major damage to exposed mobile homes. Extensive damage to poorly constructed signs. Some damage to roofing materials and buildings; some window and door damage. AND/OR: storm surge 6 to 8 feet above normal. Coastal roads and low-lying escape routes inland cut by rising water 2 to 4 hours before arrival of hurricane centre. Considerable damage to piers. Marinas flooded. Small craft in unprotected anchorages torn from moorings. Evacuation of some shoreline residences and low-lying island areas required.
3	111–130 mph (96–113 kt or 178–209 km/hr)	Foliage torn from trees; large trees blown down. Practically all poorly constructed signs blown down. Some structural damage to roofing materials of buildings; some window and door damage. Some structural damage to small buildings. Mobile homes destroyed. AND/OR: storm surge 9 to 12 feet above normal. Serious flooding at coast and many smaller structures near coast destroyed; larger structures near coast damaged by battering waves and floating debris. Low-lying escape routes inland cut by rising water 3 to 5 hours before hurricane centre arrives. Flat terrain 5 feet or less above sea level flooded 8 miles inland or more. Evacuation of low-lying residences within several blocks of shoreline possibly required.
4	131–155 mph (114–135 kt or 210–249 km/hr)	Shrubs and trees blown down; all signs down. Extensive damage to roofing materials, windows and doors. Complete failure of roofs on many small residences. Complete destruction of mobile homes. AND/OR: storm surge 13 to 18 feet above normal. Flat terrain 10 feet or less above sea level flooded inland as far as 6 miles. Major damage to lower floors of structures near shore due to flooding and battering by waves and floating debris. Low-lying escape routes inland cut by rising water 3 to 5 hours before hurricane centre arrives. Major erosion of beaches. Massive evacuation of all residences within 500 yards of shore possibly required, and of single-storey residences on low ground within 2 miles of shore.

Hurricane category	Sustained wind speed	Effects
5	> 155 mph (135 kt or 249 km/hr)	Shrubs and trees blown down; considerable damage to roofs of buildings; all signs down. Very severe and extensive damage to windows and doors. Complete failure of roofs on many residences and industrial buildings. Extensive shattering of glass in windows and doors. Some complete building failures. Small buildings overturned or blown away. Complete destruction of mobile homes. AND/OR: storm surge greater than 18 feet above normal. Major damage to lower floors of all structures less than 15 feet above sea level within 500 yards of shore. Low-lying escape routes inland cut by rising water 3 to 5 hours before hurricane centre arrives. Massive evacuation of residential areas on low ground within 5 to 10 miles of shore possibly required.

The maximum sustained wind speed used in the scale is based on the peak 1-minute wind at the height of 10 m (22 ft).

Source: NOAA

Scale; minor changes to the description of wind-caused damages are expected as the new scale is being refined. The new scale represents a move away from describing the effects of the landfall of a hurricane of a certain category, towards relying on sustained wind speed as the primary determinant. Any effect of the expected minor adjustments to the description of wind-caused damages in the NOAA 2009 Saffir–Simpson Hurricane Wind Scale are likely to be negligible from the point of view of sponsors of and investors in insurance-linked securities.

It is noteworthy that there is no Category 6 in the Saffir–Simpson scale since Category 5 is unbounded. A super-hurricane is not an impossibility, and wind speeds can exceed 200 mph. One of the main reasons the scale stops at Category 5 is that the damage at landfall is truly catastrophic, and there would be little difference between Category 5 and a hypothetical Category 6. The correctness of this logic is open to debate.

HISTORICAL FREQUENCY OF HURRICANES THREATENING THE US

Lisa: Dad! I think a hurricane’s coming!

Homer: Oh, Lisa! There’s no record of a hurricane ever hitting Springfield.

Lisa: Yes, but the records only go back to 1978, when the Hall of Records was mysteriously blown away!

The Simpsons

For rare events, samples of observed values tend to be very small, leading to a considerable degree of uncertainty in estimating their probability of occurrence. Major hurricanes certainly fall in the category of such events. Figure

Table 4.4 NOAA 2009 Saffir–Simpson Hurricane Wind Scale (currently considered experimental)

Hurricane category	Sustained wind speed	Effects
1	74–95 mph (64–82 kt or 119–153 km/hr)	<i>Damaging winds are expected.</i> Some damage to building structures could occur, primarily to unanchored mobile homes (mainly pre-1994 construction). Some damage is likely to poorly constructed signs. Loose outdoor items will become projectiles, causing additional damage. Persons struck by windborne debris risk injury and possible death. Numerous large branches of healthy trees will snap. Some trees will be uprooted, especially where the ground is saturated. Many areas will experience power outages with some downed power poles.
2	96–110 mph (83–95 kt or 154–177 km/hr)	<i>Very strong winds will produce widespread damage.</i> Some roofing material, door and window damage of buildings will occur. Considerable damage to mobile homes (mainly pre-1994 construction) and poorly constructed signs is likely. A number of glass windows in high-rise buildings will be dislodged and become airborne. Loose outdoor items will become projectiles, causing additional damage. Persons struck by windborne debris risk injury and possible death. Numerous large branches will break. Many trees will be uprooted or snapped. Extensive damage to power lines and poles will likely result in widespread power outages that could last a few to several days.
3	111–130 mph (96–113 kt or 178–209 km/hr)	<i>Dangerous winds will cause extensive damage.</i> Some structural damage to houses and buildings will occur with a minor amount of wall failures. Mobile homes (mainly pre-1994 construction) and poorly constructed signs are destroyed. Many windows in high-rise buildings will be dislodged and become airborne. Persons struck by windborne debris risk injury and possible death. Many trees will be snapped or uprooted and block numerous roads. Near-total power loss is expected with outages that could last from several days to weeks.
4	131–155 mph (114–135 kt or 210–249 km/hr)	<i>Extremely dangerous winds causing devastating damage are expected.</i> Some wall failures with some complete roof structure failures on houses will occur. All signs are blown down. Complete destruction of mobile homes (primarily pre-1994 construction). Extensive damage to doors and windows is likely. Numerous windows in high-rise buildings will be dislodged and become airborne. Windborne debris will cause extensive damage and persons struck by the wind-blown debris will be injured or killed. Most trees will be snapped or uprooted. Fallen trees could cut off residential areas for days to weeks. Electricity will be unavailable for weeks after the hurricane passes.

Hurricane category	Sustained wind speed	Effects
5	> 155 mph (135 kt or 249 km/hr)	<i>Catastrophic damage is expected.</i> Complete roof failure on many residences and industrial buildings will occur. Some complete building failures with small buildings blown over or away are likely. All signs blown down. Complete destruction of mobile homes (built in any year). Severe and extensive window and door damage will occur. Nearly all windows in high-rise buildings will be dislodged and become airborne. Severe injury or death is likely for persons struck by wind-blown debris. Nearly all trees will be snapped or uprooted and power poles downed. Fallen trees and power poles will isolate residential areas. Power outages will last for weeks to possibly months.

Source: NOAA (Landsea *et al*)


4.10 illustrates historical frequency of the North Atlantic (NA) and Eastern North Pacific (ENP) named storms, hurricanes and major hurricanes. The data includes all such storm systems and not only those that resulted in landfalls.

Climate changes affect the frequency and severity of hurricanes; the majority of the scientific community holds the opinion that the current probability of major hurricanes in this part of the world, in particular in the North Atlantic, is greater than indicated by historical averages in the observation period, and may be growing. This topic, tied to the subject of global warming, is covered later in this and in other chapters. It is important to point out, however, that we do not need to believe in global warming to see climate changes that can have an effect on hurricane activity. There is some disagreement about whether the climate changes affect both the frequency and the severity of hurricanes, and, if they do, whether they affect them to the same degree.

It can be seen that few of the tropical storms become hurricanes, and even fewer develop into major hurricanes. Landfalls are even rarer, but when they happen the results can be devastating. From the point of view of insurance-linked securities analysis, it is the probability of landfall and the subsequent damage that characterise the risk. (In rare cases, insurance-linked securities can be exposed to hurricane risk even if the hurricanes do not make a landfall. An example would be damage to offshore oil platforms. Still, the risk-exposed areas are likely to be located very close to shoreline.)

Figure 4.11 shows tracks of observed North Atlantic and Eastern North Pacific hurricanes. Only major hurricanes (Category 3 and greater on the Saffir–Simpson hurricane scale) are shown; tracks and geographical distrib-

Figure 4.10 Number of named North Atlantic and Eastern North Pacific storms by year



The lower bars represent Category 3 and greater hurricanes and the hatched bars all other hurricanes, while the top bar shows all named storm systems.

Source: National Oceanic and Atmospheric Administration

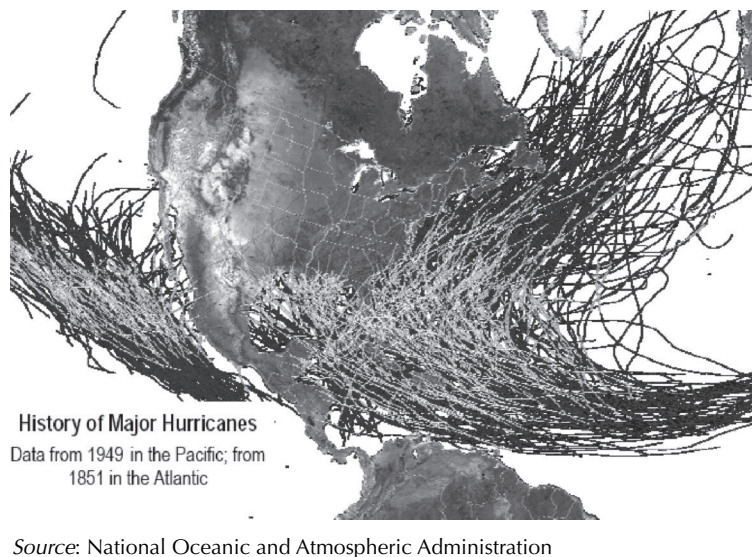
ution of formation differ by hurricane category. Florida and Texas are the two states with the greatest historical number of hurricane landfalls and damages. Hurricane risk in these two states is significantly higher than elsewhere in the coastal US. Figure 4.11 clearly shows the very high probability of major hurricane landfalls in Mexico and the Caribbean. While the majority of ILS hurricane risk in the Americas is in the US, some securities have transferred to the capital markets hurricane risk of other countries in the region, of which Mexico is the best example.

It has been suggested that the tracks have been, on average, shifting over the decades of observation. If true, this fact may be very important in probabilistic assessment of future hurricanes and their landfall locations. Unfortunately, the data is too limited to be statistically credible, and no solid argument can be made based purely on the observations of historical hurricane tracks.

SEASONALITY OF THE HURRICANE RISK IN INSURANCE-LINKED SECURITIES

The main hurricane risk of insurance-linked securities, that of North Atlantic hurricanes, is seasonal as opposed to following uniform distribution. The hurricane season officially starts on June 1 and ends November 30. Very few hurricanes occur outside the hurricane season. Approximately

Figure 4.11 Tracks of known North Atlantic (NA) and Eastern North Pacific (ENP) major hurricanes (Category 3 and greater)




97% of all tropical storm activity happens during these six months. As shown in Figure 4.12, there is a pronounced peak of activity within the hurricane season, which lasts from August through October. Over three-quarters of storms occur during this period. The percentage of hurricanes, in particular major hurricanes, is even greater: more than 95% of major hurricane (Category 3 and greater) days fall from August through October.

Definition of hurricane season is rarely used in the offering documents for insurance-linked securities. Instead, specific dates determine the coverage period. Knowing when the hurricane season officially starts and ends is not relevant. However, there are some insurance-linked securities for which the definition of the hurricane season is important. Exchange-traded IFEX catastrophe futures use a formal legal definition of North Atlantic hurricane season. This definition is used in establishing maintenance margin levels for IFEX contracts. Catastrophe futures and similar insurance-linked securities are described in detail in other chapters.

Hurricanes threatening the Pacific coast of the US and Mexico have a longer period of heightened activity, which starts earlier than on the Atlantic coast but has the same activity peak as the North Atlantic hurricanes. West Pacific hurricanes are distributed even more evenly over the year; they are less important in securitisation of insurance risk.

Figure 4.12 Distribution of hurricanes and tropical storms by month in the North Atlantic



Source: National Oceanic and Atmospheric Administration

Hurricanes in the Southern Hemisphere (called typhoons or cyclones there) tend to occur between October and May, but specific frequency distributions depend on ocean basin.


LANDFALL FREQUENCY IN PEAK REGIONS

Returning to the North Atlantic hurricanes, which present the greatest threat in the southeastern US, Figure 4.13 and Figure 4.14 illustrate hurricane landfall frequencies expressed as return periods. Unlike the figures above, only landfalls – which typically are the only hurricane risk in insurance-linked securities – are shown, with the two graphs corresponding to hurricane Categories 1 and 5 on the Saffir–Simpson hurricane scale.

Return period is defined here as the long-term average of a recurrence interval of hurricane landfalls of specific or greater intensity (category) at the time of landfall. It can also be seen as the inverse of the annual exceedance probability. Return period is usually measured in years.


Historical data is the best indicator of future hurricane frequencies. Of course, this does not mean that a simple sampling of the historical frequencies should be used in hurricane simulations. It means only that historical data is the starting point of any model, which is also where we return to validate the model once it has been built. A sound model is much more than just fitting of a distribution to the existing data points; some extremely sophisticated models have been created in recent years.

Figure 4.13 Hurricane return periods for South Eastern US: Category 1 hurricanes



Source: National Hurricane Center, NOAA

Figure 4.14 Hurricane return periods for South Eastern US: Category 5 hurricanes



Source: National Hurricane Center, NOAA

HURRICANE FREQUENCY AND SEVERITY EFFECTS OVER VARIOUS TIME HORIZONS

Continuing to focus primarily on hurricanes affecting the US, three primary phenomena affect hurricane frequency and severity, each operating over its own time scale: short term, medium term and medium to long term.

1 Short term

ENSO, which stands for El Niño Southern Oscillation, is the cycle of consistent and strong changes in sea surface temperature, air pressure and winds in the tropical Pacific Ocean. The two phases, El Niño and La Niña, typically take three to five years to complete the cycle. El Niño is the warm phase of the cycle, when the sea surface temperature in the tropical Pacific is above average. Its opposite, La Niña, is the phase when the temperatures are below average. The warming and cooling affect the level and patterns of tropical rainfall, which in turn has an effect on worldwide weather patterns and hurricane frequency and severity.

El Niño is associated with lower-than-average tropical storm and hurricane activity in the Northern Atlantic due to higher-than-average vertical wind shear resulting from the wind patterns during this phase of ENSO. The probability of hurricanes and hurricane landfalls in the Caribbean and other parts of the North Atlantic is significantly reduced during the regular hurricane season. At the same time, the weather patterns lead to an increase in tropical storms and hurricanes in the eastern tropical North Pacific. Results of the La Niña phenomenon are the opposite: storm formation and hurricane activity are increased in the North Atlantic during the hurricane season, while in the Pacific the probability of hurricanes is lower than average. These two phases of ENSO are not equal in time. El Niño rarely lasts longer than one year, while La Niña tends to take between one and three years. There is no strict cyclicity here, in the sense that each of the two phases can have shorter or much longer durations than expected. The general relationship, however, usually holds, with periods of increased hurricane activity in the Atlantic being longer than periods of decreased activity.

Technically speaking, El Niño and La Niña are not truly two phases of the ENSO cycle. The end of El Niño leads to an ENSO-neutral period, which may not be followed by a pronounced La Niña phenomenon and can instead go back to the El Niño stage. Similarly, La Niña may not be followed by a pronounced El Niño stage.

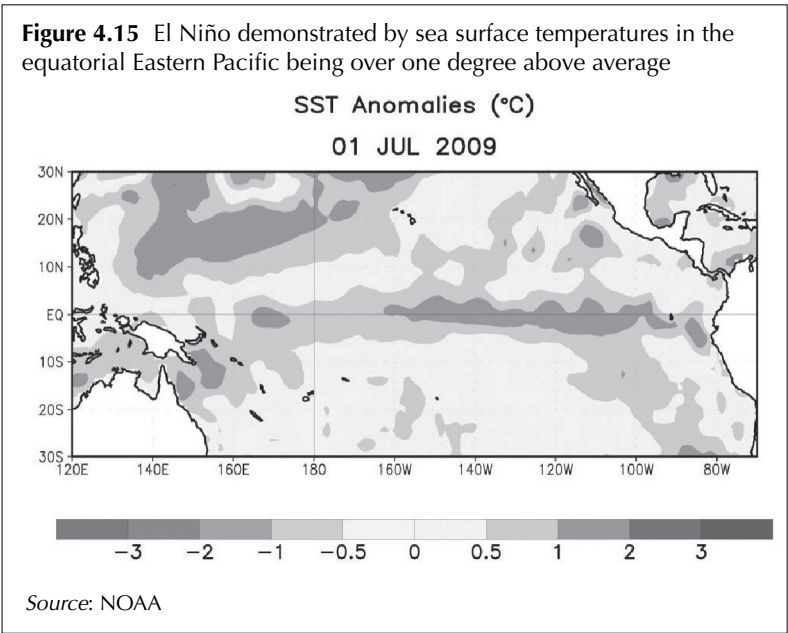
ENSO affects not only the frequency but also the severity of hurricanes. One reason for this is the vertical wind shear effect, where hurricane inten-

sity in the Atlantic is dampened during El Niño and increased during La Niña. In addition, the tropical storm formation centres differ slightly and the hurricanes follow different tracks. La Niña results not only in a greater frequency of hurricanes in the Atlantic but also in a greater probability of hurricanes being formed off the west coast of Africa. These hurricanes have a higher chance of increasing in intensity and making a landfall in the US or Caribbean as major hurricanes.

Figure 4.15 shows an anomalous increase in sea surface temperature indicative of the arrival of El Niño and the expectation of lower hurricane activity in the Atlantic.

2 Medium term

AMO, which stands for Atlantic Multidecadal Oscillation, is a cycle of consistent and strong changes in sea surface temperature in the North Atlantic. The cycle is believed to be on the order of 70 years, with the up and down phases approximately equal in time. The amplitude of the temperature variations due to the AMO is much milder than that resulting from ENSO, and the changes much slower. It is believed that we are currently in the middle of the warm phase. This phase is expected to end between 2015 and 2040.



AMO has some effect on the overall frequency of tropical storms and hurricanes, with warmer temperatures contributing to the tropical storm system development and colder temperatures leading to a reduction in tropical storms. This correlation is not strong and the effect is usually disregarded. However, during the warm phases of the cycle there is a greater chance of major hurricanes compared with the average; the chance is lower during the cold phases. This effect is unambiguous and the correlation is strong.

3 Medium to long term

Climate change, in particular the increase in seawater temperature, has a strong potential to increase both the frequency and the severity of the hurricanes landfalling on the Atlantic coast of the US. Some of the change is the result of human activities. Global warming, recognised by the majority of the scientific community, is part of the overall climate change. There is no consensus on the exact manifestations of and the speed at which climate change is happening. Some would argue that categorising climate change as having medium- to long-term effect is wrong, and that substantial changes are already happening rapidly and will accelerate. The risk of abrupt climate change triggered by concurrent development of several factors has been repeatedly pointed out. Even those who subscribe to the global-warming view without any reservations are unclear on the long-term effects of this process. In fact, some research has suggested that the increase in the seawater temperature will lead to a significant increase in hurricane activity in the North Atlantic, but that at some point the process will reverse itself and the hurricane frequency will actually decrease even if the temperature continues to rise. This, however, is a minority opinion.

While global warming remains a controversial topic, in particular because different people seem to attribute different meanings to the term, it is widely accepted that seawater temperature has been rising and that the probability of hurricanes in the North Atlantic is increasing as a consequence. This correlation has direct applications for hurricane modelling.

INVESTOR VIEWS ON MACRO-SCALE FREQUENCY AND SEVERITY EFFECTS

In the analysis of catastrophe insurance-linked securities tied to the risk of hurricanes, investors have a short-term view due to the relatively short tenor of these securities. Whether the probability of hurricanes will be greater in 15 years is not germane to the probabilistic analysis of cashflows from a

catastrophe bond that matures in two years. To the degree that long-term phenomena such as climate change are already affecting the probability of hurricanes, they are relevant to and should be incorporated in the analysis. The difficulty is in having to work with very limited data samples, because, sometimes, these can provide only anecdotal evidence of the degree to which long-term processes are already affecting hurricane development and will continue to do so within the period an insurance-linked security is expected to remain outstanding. In practice, it is currently very difficult to separate and then separately model effects of the general climate change.

Shorter-term effects such as ENSO, on the other hand, can be better modelled and incorporated in the analysis. To a lesser degree, the same is true in regard to AMO. Other processes, such as the overall warming related to climate change, are often incorporated indirectly through their influence on the observed parameters of the better-understood processes of storm formation and development.

There is a broad issue of whether, and to what degree, catastrophe models should reflect the observed increase in hurricane activity in the North Atlantic. Following Katrina and the 2004–2005 hurricane seasons in general, there was an almost universal conviction that the frequency of hurricanes in the widely used commercial models was significantly understated. (There were also concerns about how other modules of the models performed, and whether the damage and loss severity were understated.) Since then, the models have been modified to produce loss results that are greater than would be expected based purely on long-term historical data, either as the main output or as an option available to the user. The change reflects the view that the long-term observations do not represent the current atmospheric conditions that affect formation, development and landfalling of tropical storms and hurricanes. This important practical issue is discussed further below and in other chapters.

Incorporating short-term effects such as ENSO in both the models and the general analytical approach can better capture the risk profile of insurance-linked securities and provide competitive advantage to investors able to do it. For example, if El Niño starts, which can happen fast and unexpectedly, short-term probabilities of North Atlantic hurricane losses will immediately be affected. This affects the risk profile of the insurance-linked securities exposed to this risk. The knowledge of lower expected hurricane activity has immediate application in pricing new insurance-linked securities and those that can be traded in the secondary markets. Another practical application is reassessing portfolio risk and return profile in light of the information on El

Niño's start. This reassessment might identify a change in the risk and return profile of the overall ILS portfolio. The practical result would be a conclusion regarding which risk buckets have to be filled and which reduced, and the right prices for doing so.

Knowledge of expected changes in hurricane activity in the short term, along with the ability to quantify the degree of the change, can create a competitive advantage in the environment when many investors are not using proper models at all and few are able to incorporate new information in their modelling process. With some exceptions, quantifying the impact of new information such as the start of El Niño is not performed by the modelling firms. Users of the models might have a view on the adjustments to parameters that have to be made, but are unlikely to be able to properly incorporate these changes in the standard modelling tools. This area is ripe for improvement; new approaches are expected to be developed in the near future. For now, some use adjustments made primarily on judgement. These adjustments might or might not be implemented at the assumptions level, as opposed to modifying the results of modelling.

The ability to reflect short-term frequency and severity effects of atmospheric processes to properly assess risk is an advantage in trading catastrophe bonds; it is an even greater advantage in investing in and trading shorter-term instruments such as ILWs and catastrophe derivatives. There is also a question of making better predictions of landfall probabilities and associated losses of tropical storms that have already formed, which is important in "live cat" trading; but these very short-term predictions have a low degree of dependence on the macro-scale hurricane frequency effects described here.

The discussion about reflecting macro-scale frequency effects in quantifying the natural catastrophe risk in insurance-linked securities is irrelevant to most investors, since they do not attempt to make any adjustments. Their analysis might still capture some of these effects to the degree that the standard modelling software packages used in catastrophe modelling might give greater weight to recent years, as opposed to being calibrated based simply on the long-term historical record of observations. While this approach on the part of investors is inadequate and easy to criticise, it reflects the degree of difficulty of determining and quantifying the effects of macro-scale atmospheric processes on hurricane activity. A high level of expertise is required to do it properly, and there is a significant degree of uncertainty associated with these adjustments.

EVOLUTION OF INVESTOR VIEWS ON CATASTROPHE MODELLING

Incorporating short-term effects in catastrophe modelling has grown in importance over time. Given that, for catastrophe bonds, buy-and-hold used to be the only investment strategy, modelling was often performed only once. Investors rarely tried to perform any real modelling and relied fully on the analytical data in the offering circulars. Many did not do even that and based their investment decisions on other considerations, of which bond ratings were the most important. Of course, even then there were investors with deep understanding of insurance-linked securities; however, they tended to be an exception rather than the rule. Even investors with a high level of expertise in catastrophe risk, such as reinsurance companies, often based the decisions on only a rudimentary overview of the summary analysis provided in the offering circulars. Some attempts to revisit the original analysis would sometimes take place in the context of portfolio construction, with a single focus on avoiding excessive risk accumulation in some combinations of geographies and perils. Again, this statement is not universally applicable, since from the very beginning some of the players in the ILS market have been very sophisticated.

As the market has continued to develop, the level of sophistication of many investors has grown with it, even though a significant disparity remains. There are some ILS investors who lack any analytical expertise, and some who believe they understand the analytics while in reality they do not.

In general, however, the current landscape is very different from what it was in the beginning of the cat bond market. There are more new issues and bonds outstanding. There is a sizable and growing secondary market for catastrophe bonds. This creates new opportunities for portfolio rebalancing and optimisation. In addition, the ILW market has grown significantly. Catastrophe derivative markets have reappeared and are growing as well. Investors able and willing to take part in these markets and not be confined to investing in catastrophe bonds have new options to generate higher risk-adjusted return by investing in catastrophe risk insurance-linked securities. Direct hedging can be done in managing an ILS portfolio. The markets remain inefficient and liquidity insufficient, but the array of options available to investors has certainly expanded.

The ability to better model the risk has always been important in the analysis of individual securities. The better tools now available for this modelling have given investors a greater degree of confidence in the analysis and opened new options not available several years ago.

An even more important development stemming from the advances in

PANEL 4.1 RELATIONSHIP BETWEEN ILS INVESTOR SOPHISTICATION AND THE LEVEL OF ILS ANALYTICAL EXPERTISE

There is an obvious connection between the level of investor sophistication and the ability to analyse the securities being invested. However, investing in insurance-linked securities without being able to fully analyse them does not necessarily put an investor in the “naïve” category. There could be very good reasons for arriving at a well-thought-out decision not to expand resources on developing internal expertise in insurance-linked securities, but instead to allocate a small percentage of the overall funds to this asset class without performing in-depth analysis. One of the reasons could be the diversifier role that insurance-linked securities can play in a portfolio. Given a very small percentage allocation to ILS, for some investors the cost-benefit analysis might not justify developing an expertise in this asset class, though they may still have sufficient reasons for investing in ILS.

modelling catastrophic events is the ability to better model and optimise portfolios of catastrophe insurance-linked securities. The new options available to investors – more new issuances; the development of secondary markets in catastrophe bonds, combined with a greater number of outstanding bonds; the availability of ILWs and catastrophe derivatives, both exchange-traded and over-the-counter – have also increased the need for models that can be used in portfolio and risk management. The shift from the buy-and-hold investment strategy as the only available option to the ability, no matter how limited, to optimise and actively manage a portfolio of insurance-linked securities is a sea change for a sophisticated investor. Modelling insurance-linked securities on a portfolio basis has increased the emphasis on modelling. Some of the new modelling tools developed specifically for investors are described later in this chapter.

A sophisticated investor can also take advantage of the live cat trading opportunities arising when a hurricane has already formed and is threatening an area that has significant insurance exposure. Short-term forecasts can then be combined with broader portfolio modelling to take advantage of the opportunities to take on risk at attractive prices, or to offload excess risk in the portfolio. So far, very little live cat trading has been done, but at least some growth in this area is expected.

Improvement in the ability to model catastrophe risk contributes to the development of the ILS markets. Enhanced tools give investors a higher degree of confidence and open up new options. At this point, however, most

investors do not utilise the tools already available, and many make their investment decisions based primarily on judgement and a back-of-the-envelope type of analysis. While there are some extremely sophisticated players in this market, there is significant room for improvement in investor understanding and modelling of catastrophe insurance-linked securities.

ELEMENTS OF HURRICANE MODELLING

Doubt is not a pleasant condition, but certainty is absurd.

Voltaire

There is a very high degree of uncertainty associated with hurricane losses. It surrounds all elements of a hurricane model – from the frequency and location of storm formation to its tracks and intensity, and the possible land-fall and resulting insured losses. The very high degree of uncertainty has been a continuing source of frustration for many investors who rely on the output of black-box-type modelling tools such as the analysis summarised in offering circulars for cat bonds. It is even more frustrating for those few investors for whom the modelling tools are not black boxes and who understand the assumptions and the modelling of individual processes within the broader analytical framework. Their superior understanding does not eliminate the uncertainty and might even increase the perception of the degree of uncertainty in their minds. We need to keep in mind that the obvious uncertainty involved is not unique to insurance-linked securities tied to catastrophe risk: to some degree it is present in any security and financial instrument. Insurance-linked securities are unique in the type of risks they carry; they are not unique in the carrying of risk *per se*. Every security carries some degree of risk, uncertainty and unpredictability; assuming the risk is what investors are paid for. In the case of insurance-linked securities, one of the ways to reduce the uncertainty is to improve the modelling of hurricanes and the damage they cause.

There exists a considerable body of research on modelling atmospheric phenomena such as storms and hurricanes. Catastrophe models used in the insurance industry and in the analysis of insurance-linked securities are based on some of this research, as described earlier. A comprehensive overview of the atmospheric science on which the commercial models are based would take up a thick volume and cannot be provided here. In most cases, understanding all of the science is completely unnecessary for an investor analysing insurance-linked securities. It is important, however, to have some basic understanding of the science and assumptions used in catastrophe software packages and avoid treating these tools as black boxes that

spit out results based on user input. Among the many advantages of understanding the basics of the science and assumptions used by the models is the ability to better understand the sensitivity of results and the degree of uncertainty involved. Another important advantage is understanding some of the differences between the models.

Some elements of the modelling of hurricane risk and related basic scientific concepts are discussed below. They are not intended to educate a reader on the hurricane science as such, or even its use in commercial catastrophe models: rather, the purpose is to provide an illustration of how the models work, by describing selected issues relevant to the topic.

Modelling hurricane frequency

The number of storms in a hurricane season can be simulated by sampling from the hurricane frequency distribution. When the frequency of hurricanes or hurricane landfalls is modelled directly, there are three main choices for the probability distribution:

- ☐ Poisson;
- ☐ negative binomial; and
- ☐ binomial.

Poisson distribution is the natural first choice as it is for most frequency distributions. Binomial distribution might be appropriate where the sample variance is less than the sample mean. This is unlikely to be the case in events with such a high degree of uncertainty as hurricanes; the fact that there can be several hurricanes during the same time period further complicates the use of this distribution. In fact, the variance generally exceeds the mean, leading to the recent adoption by many of the negative binomial as the distribution of choice for hurricane frequency. Most of the standard catastrophe models utilise the negative binomial distribution for hurricane frequency in Florida; some allow users the choice between Poisson and negative binomial distributions.

Despite the recent shift towards the use of the negative binomial distribution, Poisson distribution is still commonly used as well. When considering the choice of probability distribution for hurricane frequency, parameterisation might be a bigger issue than the analytical form of the distribution. This is particularly challenging because of the varying views on the changes in hurricane frequencies over time. In fact, the regime switch view of the hurricane frequency affects both the choice of the parameters of the distribution and the choice of the distribution itself. It is possible that the statistically

significant fact of the sample variance exceeding the sample mean is the result of inappropriately combining in the same sample unadjusted observations from time periods that have had different mean hurricane frequencies due to climate oscillations or other changes. If this is the case, the choice of Poisson distribution over the negative binomial might be preferable. In this context, the choice of the distribution is dependent on the choice of the distribution mean: if it is determined based on the full historical database of observations, with all observations given the same weight, negative binomial distribution seems to almost always outperform Poisson in back-testing regardless of the geographical region being considered.

Hurricane frequency and intraseasonal correlation

There is an ongoing debate about whether the occurrence of a hurricane, in particular a major hurricane, during the hurricane season means that there is a greater probability of another hurricane occurring in the remainder of the season. In other words, there is a question of whether the frequency distribution changes if it is conditioned on an occurrence of a hurricane. The phenomenon in question is sometimes referred to as hurricane clustering.

The rationale for the view that the probability of hurricanes increases under these circumstances is that a major hurricane is more likely to develop if the general atmospheric conditions are more conducive than average to hurricane formation. This in turn implies a greater-than-otherwise-expected chance of additional hurricanes during the season.

In the analysis of insurance-linked securities, the issue of intraseasonal correlation is of particular importance for second-event bonds and second-event catastrophe derivatives. Of course, it is important in ILS analysis in general for valuation purposes as well as for evaluating opportunities in the catastrophe bond secondary markets. It could be of even greater consequence in the context of investment portfolio management. If the probability of hurricane losses on the US Atlantic coast has increased, it could affect several securities and have a magnified effect across the portfolio.

In practice, we would be hard pressed to find investors who go through the process of calculating conditional probabilities of hurricane events. The standard commercial catastrophe models do not have an easy way to adjust the probabilities in the middle of a hurricane season based on the occurrence of an event such as Category 3 hurricane making a landfall in the US or the Caribbean. There have been attempts to take the intraseasonal autocorrelation into account in modelling second-event catastrophe bonds. A better approach than autocorrelation models or making adjustments to the

frequency distribution based largely on judgement would be to instead adjust the atmospheric parameters in the model. If the occurrence of a hurricane was indicative of changing atmospheric conditions, then the best way to reflect it in the model is by making changes to these assumptions. The approaches of using autocorrelation methods or of making adjustments based primarily on judgement are also important.

Wind field modelling

Storm track modelling and modelling of the characteristics of the storm are an essential part of the overall hurricane modelling. Characteristics of the storm at a particular location include central pressure, direction, forward velocity, maximum winds, air pressure profile and many others.

Some elements of wind field modelling are shown in Panel 4.2. The approach shown is just one of many ways to build wind field models.

The important output of wind field models that is used in insurance catastrophe-modelling software packages is the wind characteristics after hurricane landfall, at specific locations where insured exposure is located.

Parameterisation of the models is a challenging task that has the potential to introduce uncertainty and, in some cases, lead to significant errors. While historical observations are used to calibrate and validate the models, the sample of observed events is not big enough to credibly estimate a large number of parameters. A very complex and scientifically sound theoretical wind field model might be completely useless in practice if it requires estimating a large number of parameters based on empirical data. This statement is not limited to wind field models and is applicable to most elements of hurricane modelling.

Probability distributions of some wind field parameters

In the same way as there are several wind field models, there is more than one way to model individual parameters used in these wind field models. Most wind field models use the same general parameters. Below we look at the examples of probability distributions of some of the stochastic parameters, in particular the ones used in the standard commercial catastrophe models, as these are of most interest to the practitioner.

Annual frequency

Generating storm formation frequency technically is not part of wind field modelling and comes before it, as does generating hurricane landfall frequency in most models. Hurricane frequency has been covered above,

PANEL 4.2 ELEMENTS OF WIND FIELD MODELLING

Wind field modelling is a critical part of simulating hurricanes and resulting insurance losses. Various models have been developed; even for the same model, parameterisation differs from one modeller to another. For illustrative purposes, below we show selected elements of one of the wind field models.

Pressure isobars of a cyclone can be modelled as concentric circles around its centre. One of the standard models for the radial distribution of surface pressure is

$$p(R) = p_0 + \Delta p \exp \left[- \left(\frac{R_{\max}}{R} \right)^B \right]$$

where $p(R)$ is the pressure at a distance R from the centre of the cyclone, p_0 is central pressure, R_{\max} is radius to maximum winds, Δp is the central pressure difference, and B is a scaling parameter reflective of pressure profile. There are a number of models for the Holland parameter B , one of the simplest being $B = a + b\Delta p + cR_{\max}$, where a , b and c are constant. In this formulation, dependence on latitude is taken into account indirectly through other parameters. A popular wind field simulation model is based on the gradient balance equation of the following form:

$$V_g^2 = \frac{R}{\rho} \frac{\partial p(R)}{\partial R} + (V_T \sin \alpha - fR) V_g$$

V_g is the gradient wind speed at distance R from the centre and angle α from the cyclone translational direction to the site (clockwise considered positive), ρ is the air density, f is the Coriolis parameter and V_T is the cyclone translational speed.

Using the pressure distribution model described above, we obtain the following formula for gradient wind speed:

$$V_g = \frac{V_T \sin \alpha - fR}{2} + \sqrt{\left(\frac{V_T \sin \alpha - fR}{2} \right)^2 + B \frac{\Delta p}{\rho} \left(\frac{R_{\max}}{R} \right)^B \exp \left[- \left(\frac{R_{\max}}{R} \right)^B \right]}$$

Gradient wind speed V_g can then be used to determine wind speed at various heights. A number of decay models can be used to simulate the evolution of wind parameters upon landfall. These will be utilised in calculating wind gusts over land, taking into account surface roughness and general topography.

where two functional distribution forms – Poisson and negative binomial – have been described as the most appropriate, with a general shift to using the negative binomial distribution because the variance of observed hurricane frequencies typically exceeds its mean. Parameters of the distribution, whether negative binomial or Poisson, are estimated based on a smoothing technique to account for the low number or lack of observations in most locations.

Landfall locations

If the landfall frequency is estimated directly by location based on one of the methods described above, there is no need to use any distribution to estimate landfall location probabilities. Otherwise, given the general hurricane landfall frequency, the probability of landfall by specific location can be distributed based on smoothing of empirical data or using a physical model. Other approaches can be used as well.

Central pressure

Smoothed empirical distributions can be used for central pressure at and following landfall. The same approach is possible but harder to implement for modelling hurricane central pressure before landfall. While central pressure does not easily lend itself to being described by any standard functional probability distribution, the use of Weibull distribution has produced acceptable fit. Strong hurricanes are much rarer than the weak ones, and the Weibull distribution, with properly chosen parameters, captures this relatively well.

Forward speed

Smoothed empirical distribution specific to a landfall gate is one of the choices for modelling hurricane forward speed. Similar to the central pressure distribution, that of forward speed is skewed, with very fast forward speeds being much less common than slower speeds. However, based on historical observations, the degree of skewness is generally lower. Lognormal distribution is a good choice for modelling storm forward speed in most situations.

Radius to maximum winds

Lognormal distribution can be used for modelling R_{\max} , with its parameters depending on central pressure and location latitude. The lognormal distribution needs to be truncated to avoid generating unrealistic values of R_{\max} .

Gamma distribution has also been used for stochastically generating radius to maximum winds, producing acceptable results when limited to modelling the R_{\max} variable at landfall as opposed to including its modelling over open water. Another way to generate R_{\max} values is by using one of the models where logarithm of R_{\max} is a linear function of central pressure (and/or its square) and location latitude. Coefficients in the linear relationships are determined based on empirical data. Then R_{\max} is not simulated directly, but rather is calculated as a function of latitude and the simulated value of central pressure. Other models can also be used.

These are just some of the random variables simulated in catastrophe models. Many others need to be modelled, including such important ones as wind dissipation overland, in order to ultimately derive hurricane physical parameters after a landfall.

DAMAGE MODELLING

In catastrophe models, the next step after simulating physical effects of a hurricane (such as peak gusts and flood depth at specific locations) is determining the damage they cause. Conceptually, this process is very straightforward. It involves the following basic steps:

1. For each individual location in the insured exposure database, consider
 - ☐ simulated physical characteristics of the storm that are relevant to estimating potential damage;
 - ☐ characteristics of the insured property at the location.
2. Identify the damage functions corresponding to the hurricane's physical parameters (peak gusts) and the vulnerability classes of insured buildings and contents at the location.
3. Apply the damage functions to the replacement value of the insured property to calculate the loss.

Detailed information on the insured property is essential for assessing its vulnerability to hurricanes. The information should include the following, in as great detail as possible:

- ☐ precise location of the insured property (street address, ZIP code, CRESTA, etc.);
- ☐ vulnerability characteristics (construction type, height and footprint size, year of construction, occupancy type, mitigating factors, etc.); and
- ☐ replacement property value.

Vulnerability functions are based on historical data and structural engineering analysis. Their details represent a highly proprietary component of commercial catastrophe models that can be a significant differentiator among the models. The exact definition of a vulnerability function is the relationship between the mean damage ratios and the peak gusts, where the mean damage ratio relates the expense of repairing the damaged property to the replacement cost of the property.

Modifications to vulnerability functions or subsets of vulnerability functions can be based on secondary characteristics or mitigation measures such as roof type, roof strength, roof-to-wall strength, wall-to-floor and wall-to-foundation strength, opening protection and others. The variables are largely the same for all models since they are a function of the type of exposure information collected by insurance companies. The way vulnerability functions are determined and modified differs, sometimes significantly, from one model to another. Some models use additional variables such as wind duration to better estimate damage to insured property from hurricanes.

The fact that damage modelling follows very simple and logical steps does not imply the ease of building a module for its calculation as part of a catastrophe model. The effort going into determining and refining vulnerability functions cannot be overestimated. Complex structural engineering studies have been conducted for this purpose and a large amount of historical hurricane damage data has been analysed. This is a continuing process as more precise site information becomes available, building codes change and other developments take place.

FINANCIAL LOSS MODELLING

Once the damage for each insured location has been calculated, it can then be translated into the amount of insured loss by applying to it policy terms and conditions including its deductible and limit. Loss triggers, insurance coverage sublimits and other factors are also taken into account in the calculations; for reinsurance purposes, other factors such as attachment point are also part of the loss calculations. This process too is very straightforward in its implementation as long as all the necessary data inputs are reliable.

Adjustments to the process, when such are required, can introduce a degree of complexity. Adjustments include taking into account demand surge following a catastrophic event.

PANEL 4.3 WIND AND EARTHQUAKE STRUCTURAL ENGINEERING ANALYSIS

The ability to estimate potential damage to insured structures depending on the physical characteristics of a hurricane or an earthquake is a challenging structural engineering task. Two separate disciplines, hurricane engineering and earthquake engineering, have developed to deal with engineering aspects of hurricane and earthquake hazards. While the broader focus of the disciplines is on designing, constructing and maintaining buildings and infrastructure to withstand the effects of catastrophic events, in insurance catastrophe modelling the emphasis is on quantifying the damage that would result from hurricanes and earthquakes of various intensities. Similar principles can also be applied to the risk of manmade catastrophic events such as acts of terrorism.

Estimating the dependence of mean damage ratios on hurricane peak gusts or earthquake physical characteristics for various types of structures is the process of constructing vulnerability functions, which are an essential part of the damage calculator in insurance catastrophe models. Constructing sets of vulnerability functions for specific geographical areas is necessary to take into account the overall topography, building codes and the history of their change over time, and other factors.

Demand surge

A catastrophic event such as a hurricane landfall or an earthquake can result in the increase of costs of repairing the damage and other expenses covered by insurance policies above the level of claim costs expected under normal circumstances. This effect is referred to as demand surge, reflective of the increase in costs being driven by a sharp increase in demand while the supply lags behind. An example is the shortage of building materials following a major hurricane, when many properties are damaged and all of them require building materials for restoration, all at the same time immediately following the hurricane. The cost of building materials naturally goes up to reflect the demand–supply imbalance created by catastrophic events. The post-event shortage expands to the labour costs, which also affect the cost of rebuilding the damaged property. Additional living expenses can also grow after a large catastrophic event, further contributing to losses suffered by insurance companies.

To account for demand surge, insurance catastrophe models can apply special demand surge or loss amplification factors to insurance losses. The

greater the magnitude of a catastrophic event, the greater the demand surge effect. The effect applies to different parts of insurance coverage to different degrees; consequently, demand surge factors differ as well. Sometimes the factors are further refined to reflect the various degrees of the demand surge effect, for example on the cost of rebuilding various types of property.

Aggregate approach

An aggregate approach, as opposed to the more detailed location-by-location modelling, starts before the financial loss module, in the analysis of hurricane damage. The goal here is to arrive at aggregate insured losses for an individual risk portfolio or even for the whole insurance industry. In this approach, portfolio-level information is used in the calculations to arrive at the loss distribution, as opposed to analysing each individual risk independently and then aggregating the losses across the portfolio. Inventory databases of property exposure are utilised to help accomplish this goal, with the data aggregated by location (such as ZIP or postal code) and including information on the types of property, vulnerability degrees, type of coverage, etc. The calculations consider aggregate exposure data by location, estimate the average damage and then translate it into financial losses. When this is done not for an individual portfolio of a specific insurance company but for the whole insurance industry, the result is a figure for industry-wide losses by geographic area (for example, all of Florida), the probability distribution of which is important for larger primary insurance writers, and even more important for reinsurance companies.

There are other ways to calculate aggregate losses, which are based on more granular analysis and the use of databases of insurance policies from several insurance companies, and then extrapolating the losses to the total insurance industry based on insurance premiums or another measure of exposure. Some modelling companies might have developed such databases by combining data from the companies that provided them with this information.⁴

In the context of insurance-linked securities, aggregate losses suffered by the insurance industry are important in catastrophe bonds with an industry loss trigger, in industry loss warranties (ILWs) and in catastrophe derivatives.

CATASTROPHE MODEL STRUCTURE

A catastrophe model that can be used in modelling insurance losses includes all the primary elements mentioned above. It starts with generating a

natural catastrophe event such as a hurricane or an earthquake, then determines its physical characteristics at the locations where insured properties are situated, and finally determines the degree of damage caused to the properties and the total financial loss to the insurance companies.

The model effectively simulates many (sometimes as high as a million or even more) hypothetical years and accumulates the loss statistics over these hypothetical years. The large number of simulations is essential when dealing with very rare events.

The basic structure of the catastrophe models has been described in this and the previous chapter. Figure 4.16 shows a structure of a catastrophe model that is designed specifically for the hurricane hazard; it also shows some of the parameters that are generated by the model in intermediate steps in order to arrive at the final result, aggregate financial loss.

Most (but not all) modules of the model are relatively independent of each other, with one feeding its output into the next one. Each module is critical in that it affects the end result to a significant degree. This structure explains the need for the wide-ranging multidisciplinary expertise required for developing such a model.

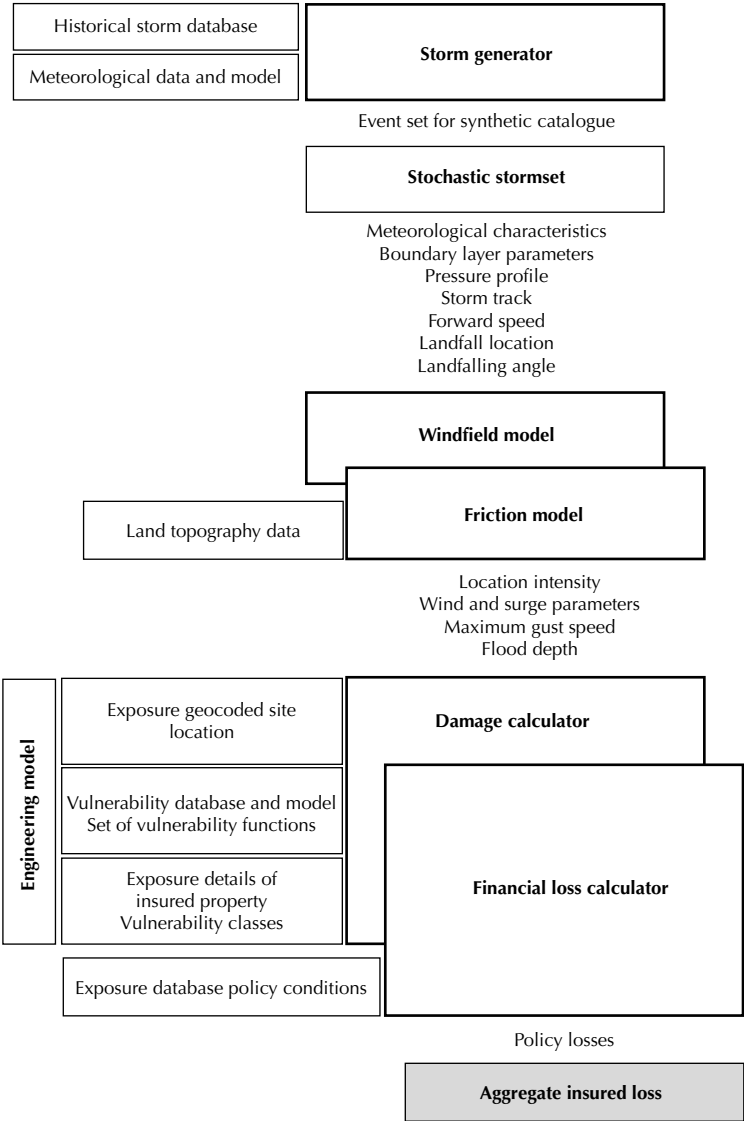
The distribution of aggregate insurance losses is the primary piece of information used in the analysis of indemnity catastrophe bonds. A model like the one outlined in Figure 4.16 also allows us to produce the probability distributions of total industry losses or of catastrophic events without referencing insurance losses, which are needed in the analysis of catastrophe bonds with industry loss and parametric triggers respectively. Not all elements of the model might need to be utilised in these cases.

MODELLING TERRORISM RISK

Modelling the risk of terrorist attacks has unique challenges not present in modelling natural catastrophes. Similar to natural catastrophes, acts of terrorism are represented by a sample of historical observations. However, the applicability of such data to the present can be limited in that the political, societal and technological landscape has probably changed since the historical observations were made. Until September 11 of 2001, our assessment of potential terrorist attacks was certainly different. In addition to the changing sociopolitical and technological landscape, there is also the human factor of terrorists dynamically trying to choose the targets, weapons and operational means of implementing an attack.

The chapter on securitising extreme mortality risk provides an overview of how the risk of terrorism was modelled in some of the extreme mortality

Figure 4.16 Hurricane catastrophe model structure



bonds. In summary, the model developed by Milliman for those transactions was based in part on a multi-level logic tree approach. At each level of the logic tree, three choices were possible: “success” of the terrorist attack, resulting in a random number of deaths in the predetermined range;

“failure” of the terrorist attack; and escalation to the next level of severity (greater number of deaths). The third choice led to the next level of the decision tree, where the same choices were presented. At every level, probabilities of each outcome – “success”, “failure” and escalation – were determined by fitting a distribution to the actual observations over the previous six-year period (that included 2001). The model was simple and based on a very limited number of observations; however, it is not clear that more mathematically sophisticated models add value unless they are based on additional external information.

The terrorism model described in the chapter on extreme mortality securitisation focuses entirely on the risk of mortality due to acts of terrorism. Property and other damage resulting from terrorism was not directly modelled.

Risk Management Solutions (RMS) has developed its own proprietary terrorism risk model for the US, as well as a global model. The model is based in part on the game theory approach to reflect changes in the landscape. The situation is constantly evolving: as antiterrorism measures and higher security are implemented, terrorists change their tactics and potential targets. The moving target creates modelling difficulties that cannot be addressed in a mathematical model but require extensive expert input. In fact, this might be one of the cases where scenario analysis is preferable to a fully probabilistic framework.

Using expert input is required to first build a database of potential targets. Prioritising the targets is the next step; it requires the analysis of both the target’s attractiveness to a terrorist and the degree of the target protection. As the latter factors change, the priorities are adjusted as well. The database of potential targets also contains data on potential damage to life and on economic loss from a terrorist attack.

A terrorism model should also incorporate the fact of the existence of several attack modes based on various weapons that could be used. In addition to conventional weapons, chemical, biological, radiological or nuclear (CBRN) weapons can be utilised, each with its own probability of occurrence and potential damage. The choice of terrorism weapons can also be site-specific, as some weapons would be more natural choices for attacks on specific sites. Finally, the mode of attack might be unconventional but it might not fit in the CBRN category either. The attack on the World Trade Center in 2001 provides an example of such a type of weapon.

The RMS probabilistic terrorism model is a bold attempt to combine rather sophisticated approaches taken from game theory, with extensive

input on potential targets, threat levels and terrorist behaviour modes, in order to quantify the risk of losses from terrorism, with the focus on large losses that can be called catastrophic. The input is dynamic in that the new developments such as antiterrorism measures, information on potential types of weapons that might be in the hands of terrorists, and even the level of “chatter” detected by the intelligence community can in theory be reflected in the inputs into the model or in adjusting some of its parameters. The overall framework appears to allow a growing degree of sophistication and the incorporation of additional information on a dynamic basis. The practical implementation, however, presents numerous challenges.

In assessing a difficult-to-quantify risk such as terrorism, it is particularly important to augment the probabilistic approach with scenario analysis. Along with allowing for reasonability testing, scenario analysis introduces one more way to use expert judgement in analysing exposure to the risk of terrorist attacks.

MODELLING PANDEMIC FLU RISK

The risk of a global pandemic of an infectious disease is not insignificant. The chances of a pandemic of a serious disease with a high level of mortality might be small, but the consequences of such an event would be catastrophic. Focusing on insurance losses, there would be a spike in mortality rates resulting in life insurance losses of possibly a catastrophic nature, as well as an avalanche of medical claims resulting in huge health insurance losses. The latter might be the case even if the mortality rate is not high but the severity of the disease is. Finally, there would be property-casualty insurance losses. These would obviously include business-interruption insurance losses. However, it is possible that other lines of property-casualty insurance business might suffer even greater losses, even though such losses are usually not fully contemplated in catastrophe risk analysis.

The chapter on extreme mortality bonds describes how pandemics have been modelled in the context of evaluating their potential impact on mortality rates resulting in a mortality spike. In analysing the risk of pandemics, the main focus is flu pandemics, since these are considered to represent the great majority of this type of risk in modern times. As described in Chapter 12, Milliman created a model for analysing the risk of mortality spikes due to flu pandemics in catastrophe mortality bonds. The model separated the frequency and severity components, parameters of which were estimated based on the available historical data. The data for frequency was considered over a long (multi-century) period of time, at least

in some cases. Binomial distribution was used for annual frequency, which is a natural choice in modelling the frequency of such events. Severity data was based on five or six data points in the more recent history. In at least one of the securitisations, Milliman modelled severity as a percentage of excess mortality fitted to these historical data points, one of which was adjusted by placing a cap on broad mortality improvements in the general population. (See the fitted severity curve for excess mortality resulting from pandemics for the Tartan Capital securitisation, in the chapter on the securitisation of extreme mortality risk.) The Milliman model then simulates the pandemic results by sampling from the frequency and severity distributions. The current Milliman model's results are sensitive to the distribution of age and gender.

The binomial frequency distribution assumes that the probability of a pandemic is the same in any year. It is likely that the current risk of a flu pandemic is elevated above the average historical levels. This can be reflected by adjusting the mean of the binomial distribution; significant judgement and expert input are required to properly make this adjustment.

The Milliman model is of the type that is sometimes called actuarial, in that frequency and severity are modelled separately based on available historical data. Another approach – the epidemiological one – is used in the model developed by RMS. It is based on a standard epidemiological approach known as SIR modelling (susceptible, infectious, recovered), which allows us to take into account additional variables such as vaccination, immunity, viral characteristics and lethality in a more direct way. The RMS model presents a more sophisticated approach from the mathematical point of view; but whether it is better than the simpler Milliman model is not fully clear, since it requires a number of inputs that introduce uncertainty and have the potential to skew the results. In the longer term, however, the RMS model is likely a better one to use for modelling pandemic risk.

The Swiss Re internal model is reported to be a combination of the actuarial and epidemiological types. The excess mortality rates are estimated based on historical data as in the Milliman model, but are then adjusted to take into account the changes that have happened since those observed events. These changes include new virus threats, vaccinations, better standards of medical care, etc. A significant degree of judgement is used in making these adjustments.

The chapter on securitisation of extreme mortality risk shows a fully stochastic model of the spread of a pandemic, implemented on the Los Alamos National Lab supercomputer. This approach is probably the one

that will eventually become the standard. Right now it is not realistic. Of the models described above, the RMS model is the closest to this approach.

PRACTICAL MODELLING OF CATASTROPHE RISK

It is not certain that everything is uncertain.

Blaise Pascal

The time of occurrence of a natural catastrophe is unpredictable. Its magnitude is unpredictable too. So is the damage it causes in its wake. This is the inherent uncertainty associated with such events as hurricanes or earthquakes. When it comes to natural catastrophes, we are in the country where predictions do not work. Manmade catastrophes are in the same territory.

The goal of modelling catastrophic events in the context of insurance securitisation as well as in general is to minimise the uncertainty surrounding the probability distribution of possible outcomes. The closest to certainty is the one who most precisely identifies and quantifies the uncertainty of these random variables.

Available models

The previous chapter identified the three main providers of commercial catastrophe-modelling software used in the analysis of potential insurance losses. In addition to AIR Worldwide, EQECAT and Risk Management Solutions, there are additional providers of either software or consulting services based on proprietary software for modelling of catastrophic insurance losses. These tend to focus on one type of hazard in a specific geographic area. For example, Applied Research Associates' hurricane model and URS's earthquake models (combined and modified under the Baseline Management umbrella) are now covering all of the US. There are also some noncommercial models such as the Florida Public Hurricane Loss model (for Florida hurricane risk only) and FEMA's HAZUS tool, which in its modified form can be used for modelling insurance losses.

While a number of external models exist, in practice only the main three, AIR Worldwide, EQECAT and Risk Management Solutions, have been utilised in securitisation of insurance risk. This is reflective of the complete domination of these three companies in the insurance and reinsurance industry and the credibility they have earned over the years. Problems – real or perceived – with modelling software developed by these companies have been pointed out on a number of occasions. However, they do have the track record and credibility that no competitor possesses.

Some companies in the industry, in particular reinsurance companies,

have developed their own proprietary models of insurance catastrophe risk. However, these are generally not full catastrophe models but rather the software that sits on top of the three established models and uses their output to obtain its own estimate, which might be different from the results of each of the underlying models.

While not every peril in every geographical area can be modelled, there now exist catastrophe models covering all the key areas of insurance exposure. Table 4.5 shows an incomplete list of the existing peril models and the countries for which they have been created. In almost all circumstances, all three major modelling companies would have these models.

While many individual models – for specific perils and countries – are available, not all of them have the same degree of credibility. Models for some regions and perils are based on more extensive research and have existed for a longer period of time. The longer period of time has created more opportunities for model validation and refinement. Not surprisingly, the three most refined models cover:

1. North Atlantic hurricanes (in particular Florida and the other Gulf states in the US);
2. California earthquakes; and
3. Japanese earthquakes.

These three represent the biggest catastrophe risks for the insurance industry. They combine high concentration of insured exposure and high probability of catastrophic events. Even though the models produced by the three modelling firms have existed for a long time, their results differ, sometimes significantly, from one firm to another, and significant adjustments to each of them have been made even very recently. The net result is the uncertainty that still exists in quantifying catastrophe insurance exposure even in the areas where the research has been extensive and the investment in model development quite sizable.

It is important to carefully analyse whether indirect effects of natural catastrophes have been modelled, and, if so, how. These indirect effects include, for example, flood following a hurricane and fire following an earthquake. These secondary effects might result in more damage than the primary ones, and their proper modelling is critical.

Unmodelled losses

One of the most common examples of unmodelled losses are those that reflect improper data coding, resulting in wrong or incomplete entry of

Table 4.5 Catastrophe model availability by peril and geographical region/country†

Peril	Region	Country†
Hurricanes, cyclones and storms	North America, Mexico and Caribbean	US (including Alaska), Mexico, Bahamas, Barbados, Bermuda, Cayman Islands, Dominican Republic, Jamaica, Puerto Rico, Trinidad and Tobago
	Europe	Austria, Belgium, Denmark, France, Germany, Ireland, Netherlands, Norway, Sweden, Switzerland, UK (including flood)
	Asia-Pacific	Australia, China (including Hong Kong), Hawaii (US), Japan, Philippines, Taiwan
Earthquakes	North America, Mexico and Caribbean	US (including Alaska), Canada, Mexico, Bahamas, Barbados, Cayman Islands, Dominican Republic, Jamaica, Puerto Rico, Trinidad and Tobago
	Central and South America	Belize, Chile, Costa Rica, Colombia, El Salvador, Guatemala, Honduras, Nicaragua, Panama, Peru, Venezuela
	Europe and Middle East	Greece, Israel, Italy, Portugal, Switzerland, Turkey
	Asia-Pacific	Australia, China, Hawaii (US), Indonesia, Japan, New Zealand, Philippines, Taiwan
Tornado and related	North America	Canada, US
Terrorism	North America	US (worldwide terrorism models also exist but their credibility level is unclear)
Flu pandemic	Worldwide	Worldwide

†List incomplete; countries with less significance to insurance securitisation might not be shown.

exposure into the model. This is part of the pervasive issue of data quality described below.

It is not unusual for some of the insured exposure not to be reflected in the models because they are not designed to handle specific types of coverage. Additional perils, related to the main one but in an indirect fashion, would probably not be taken into account by the model. Finally, there might be

insurance losses due to catastrophic events that have never been contemplated in the original coverage but still have to be paid by insurance companies. Care should be taken to make sure that all losses that can be modelled by catastrophe software are input, and any other losses evaluated separately.

Modelling results presented to investors

As a reminder of the primary goal of the analysis, Panel 4.4 shows the summary output of the risk analysis performed for an indemnity catastrophe bond (see the chapter on property catastrophe bonds for additional information). It is no more than a summary, but it is often the main part of the information included in the offering circulars, no matter how long the risk analysis section appears to be.

DATA QUALITY

The quality of data used in catastrophe models is as important as the quality of the models themselves. Data used to create and parameterise the models affects the precision and correctness of modelling results. Many elements of the existing models have been built so that they can take advantage of the most reliable data available. For example, certain hurricane data available from the National Oceanic and Atmospheric Administration databases include measurements at six-hour intervals. Models have been constructed specifically to take the six-hour intervals into account, as other data is either unavailable or not fully reliable. This is also the data used to validate the models.

The issue of data quality is usually raised not in the context of the data used to formulate and parameterise the models, but in assessing the reliability and completeness of the data on the details of the exposure in applying a catastrophe model to a portfolio of insurance policies. Quality of the insurance data serving as input into catastrophe models is an industry-wide issue introducing a significant degree of uncertainty to results of the modelling process. Best practices are still in the process of being developed, and the quality of data can vary widely from one insurance company to another. Improper data coding or not capturing all the relevant exposure data in sufficient detail is also an indication of deficiencies in the underwriting process.

Implications for investors can be significant. Two insurance-linked securities, such as catastrophe bonds with indemnity trigger, might appear very similar but in reality have different risk profiles because of the different

PANEL 4.4 ILLUSTRATIVE SUMMARY OUTPUT OF RISK ANALYSIS OF A CATASTROPHE BOND

A simplified catastrophe bond description is presented below. The coverage attaches at US\$5 billion of ultimate net loss resulting from a single occurrence of a hurricane.

Transaction parameters

Covered risk	Hurricane affecting specific insurance portfolio
Trigger	Indemnity per occurrence (UNL)
Attachment level	US\$5.0 billion
Exhaustion level	US\$5.5 billion
Insurance percentage	50%
Principal amount	US\$250 million

Based on the per-occurrence exceedance probabilities resulting from catastrophe modelling of the subject insurance portfolio, key risk measures are calculated. The expected loss in this example is 1.48% per annum. The attachment probability is 1.70%.

Risk measures	Base case (standard catalogue) (%)	Warm Sea Surface Temperature catalogue (%)
Attachment probability	1.70	2.54
Exhaustion probability	1.30	1.83
Expected loss	1.48	2.15

In this example, modelling was done twice: first with parameterisation based on the long-term historical averages of hurricane activity in the covered territory, and then based on the so-called Warm Sea Surface Temperature catalogue to take into account the greater chance of hurricane activity in the current period. The latter is of most interest since it is believed to present results that are more realistic.

This summary does not include many of the other important elements of risk analysis. However, it does show the two figures of most interest to investors: expected loss and attachment probability. Expected loss provided in the offering circular serves as the starting point for analysis performed by investors.

degrees of uncertainty related to data quality and underwriting standards in general. In evaluating such insurance-linked securities, the few investors familiar with underwriting processes of individual insurance companies can have an advantage over those not possessing this level of expertise.

The seemingly inconsequential issue of data quality can play a much greater role in modelling catastrophe risk than we would expect. It presents a good illustration of the “garbage in, garbage out” principle, and could be an important element of the analysis performed by investors.

INVESTOR AND CATASTROPHE MODELLING

Investors in catastrophe insurance-linked securities are presented with numerous choices and decisions in their analysis. Most of them have been mentioned or alluded to above.

The questions to be answered are numerous. Which catastrophe model is most appropriate for a specific type of risk exposure? How different are the results of different models? Are there known biases in some models related to specific perils or geographical regions? Are models for one region more credible than for another? How can we quantify the additional uncertainty related to the lower credibility of some models? Are there ways to validate some modelling results? What are the primary sources of uncertainty in the modelling? How do we quantify the additional uncertainty of securities with indemnity as opposed to parametric trigger?

The list of questions never ends, which once again underscores the advantages of having modelling expertise in the analysis of insurance-linked securities. It almost makes us wonder whether the informational disadvantage of the investor is too great to play the ILS game. The disadvantage is relative to both the sponsors of catastrophe bonds and to reinsurance companies that often invest in these securities. Both seem to have the level of expertise that an investor is usually unable to achieve. The answer to this question is more optimistic than it appears to be, however. Investors can and do participate in this market and generate attractive risk-adjusted returns. While reinsurance companies in their role as investors seem to have some expertise that few investors possess, it is not necessarily the type of expertise that is most important in ILS investing. Investors have the capital markets outlook that is usually lacking in insurance and reinsurance companies investing in insurance-linked securities. This capital markets view gives investors an advantage in some areas even when they are disadvantaged at others.

Ultimately, the conclusion is simple: modelling is critical, and without

modelling expertise it is impossible to generate high-risk adjusted returns on a consistent basis. The industry is slowly coming to this realisation.

Managing catastrophe risk on a portfolio basis is one of the most critical elements of ILS investing. The choice of modelling tools is now available for this purpose; it is also discussed in the chapter on modelling portfolios of catastrophe insurance-linked securities.

CATASTROPHE BOND REMODELLING

Almost every cat bond transaction has involved the analysis performed by one of the three main modelling agencies, AIR Worldwide, EQECAT and RMS. The summary of the analysis is included in the offering documents; a data file such as an Excel spreadsheet might also be provided as part of the offering circulars. This raises the question of the differences between models. The annual expected loss or probability of attachment calculated by AIR Worldwide might differ, perhaps significantly, from the annual expected loss or probability of attachment if they were calculated by one of the other models based on the same data.

Leaving aside for a moment the question of which model is “better”, in the ideal world an investor would like to see the analysis performed by all three modelling firms and then make their own conclusions. “Remodelling” refers to analysing a catastrophe bond by a modelling firm that did not perform the initial analysis that was included in the offering documents and used in pricing of the bond. If the security has a parametric trigger, all the data is available and another modelling firm can easily perform its own analysis so that the results can be compared. Comparison is much more difficult for indemnity catastrophe bonds. For these bonds, it is necessary to have full exposure information in order to perform the analysis. Such information is never provided to investors; only summaries are included in the offering circulars.

In order to perform the analysis, in this situation another modelling firm has to make a choice between two simplifying assumptions. One of them is to assume the correctness of the analysis, such as the values of expected loss, attachment probability and the exhaustion probability. Based on these figures and the exposure summary in the offering circular, the modeller then tries to work back to the inputs to arrive at exposure expressed at a greater level of detail than is provided in the documentation. The exposure information is important in portfolio management, where it allows us to monitor exposure accumulation over many securities and properly establish the risk–return tradeoffs on a portfolio basis.

Another choice would be to start with the exposure summary in the investor documents, and try to estimate what the exposure is at a more detailed level. This could be done by supplementing the exposure data provided with publicly available data on the geographic and line-of-business distribution of exposure for the sponsor, as well as the possible knowledge by the modeller of the underwriting processes of the sponsor. The resultant expected loss and the exceedance probability would then differ from those in the offering circular.

This type of analysis can now be performed very fast, even during the initial marketing stage before the bond pricing has been finalised. This topic is revisited later in greater detail.

HURRICANE FORECASTING

“Hurricane forecasting” refers to probabilistic predictions of hurricane activity in the short term. These are not actual forecasts but probability distributions of potential outcomes based on the most current data. These forecasts refer to the upcoming hurricane season or a season already in progress.

William Gray, for all intents and purposes, pioneered the field of hurricane forecasting. He developed a number of forecasting methodologies with a special focus on North Atlantic hurricanes. Phil Klotzbach, who has taken from him the leadership of the hurricane forecasting project, in 2009 started issuing 15-day forecasts in addition to the seasonal ones. This is a big change from issuing forecasts from the one to five times a year common for hurricane forecasters. The Klotzbach/Gray group has proven its skill over the years of issuing hurricane forecasts for the North Atlantic. Its methodology is continuing to evolve, but in most general terms it is based on identifying and monitoring several atmospheric and/or oceanic physical variables, either global or relatively localised, that are relatively independent of each other and have been shown, by utilising statistical analysis tools, to serve as good predictors of the following North Atlantic hurricane season.

NOAA issues hurricane forecasts too, as do several research groups around the world. It appears that as of 2009 only the Klotzbach/Gray group has been able to clearly demonstrate its skill in forecasting probability of major hurricane landfalls in the US. Other groups either do not issue forecasts associated with landfalls or have not been recognised for their skill in successfully forecasting landfalls. In insurance catastrophe modelling, landfalls are of major importance, while hurricanes that bypass land are of interest only if they have the potential to damage oil platforms.

The forecasts create additional opportunities for optimising risk-adjusted return on a portfolio basis. They also provide input into pricing of all affected insurance-linked securities, and in particular ILWs, securitised reinsurance and catastrophe bonds close to expiration.

Live cats

The term “hurricane forecasting” is also used in reference to probabilistic assessment of development of the storms and hurricanes that have already formed and might make a landfall. The ability to trade the risk of natural catastrophic events that can occur in the very near future – from several days to several hours – creates opportunities for those who can obtain better information on the projected path and potential damage from the hurricane and to better take advantage of the situation. It also creates opportunities to offload excess risk if necessary. This “live cat” trading can be done on a more intelligent basis when short-term hurricane forecasts have a relative degree of credibility.

The topic of hurricane forecasting is revisited in the chapters on ILWs and catastrophe derivatives and on managing investment portfolios of insurance catastrophe risk.

CLIMATE CHANGE

The trouble with our times is that the future is not what it used to be.

Paul Valéry

Climate change has been mentioned more than once in the context of modelling catastrophe risk. The expectations of the future climate state are different from its current one. The effects of climate change relevant to hurricane activity, in particular the increase in sea-surface temperature, can already be observed. These changes make it harder to rely on the old approach of forming conclusions about future natural catastrophe activity based entirely on prior historical observations. The future frequency and severity of hurricane events might be a function of atmospheric and oceanic processes that are different from the ones in the period of historical observations.

The focus of an investor in the analysis of insurance-linked securities tied to the risk of natural catastrophes is on the relatively short time horizon. Changes expected to take place over a long period of time are of less significance due to their minimal impact on catastrophe-linked securities that tend to have short tenor. Unless there is a clearly observable trend, this view suggests disregarding recent changes and relying primarily on the long-

term averages of hurricane frequency and severity. If the speed of the climate change is rapid, though, this view might be incorrect; there is a need also to reflect the developing new environment in evaluating the risk of future hurricanes. In addition, it is possible that the climate changes have already altered the atmospheric and oceanic processes, probably starting a number of years ago. This view would necessitate immediately taking climate change into account. In simple terms, we can then see the observed historical sample of hurricane activity as consisting of two parts: the first, longer, period when the conditions were relatively constant and the variability was due to natural statistical fluctuations; and the second period encompassing more recent years when a trend might be present in the changing atmospheric and oceanic conditions that influence hurricane activity. The trend might be accelerating, as suggested by all of the global warming theories.

The decision regarding whether we are in the period of heightened hurricane activity and whether this activity is likely to accelerate in the very near future is an important one both for insurance companies with significant hurricane risk accumulation and for investors in catastrophe insurance-linked securities. The majority have decided that we are now in a period of climate change that has higher probability of hurricane activity than suggested by long-term historical averages. The modelling firms have incorporated this approach by creating an option in their software models to allow users to make their own choice about whether to base the analysis on long-term averages or assume higher levels of hurricane activity than suggested by the history. The latter option is referred to as using the Warm Sea Temperature Conditioned Catalogue of events when no additional trends are taken into account.

The decision to use higher levels of potential hurricane activity as the primary modelling approach is not tied directly to the acceptance of the global warming theory; as mentioned earlier, the shorter-term climate processes of an oscillating nature can provide a sufficient reason for believing we are in an environment more conducive to hurricane development than in the past.

SPONSOR PERSPECTIVE ON MODELLING

The importance of catastrophe modelling for insurance and reinsurance companies is apparent. Modelling catastrophe insurance risk is part of the enterprise risk management (ERM) process. Its results are used in making decisions on the best ways to employ company capital. They are an impor-

tant input in decisions on whether to retain the risk, reinsure some of it or transfer it to the capital markets. The transfer to the capital markets can be in the form of sponsoring insurance-linked securities such as catastrophe bonds or in the form of hedging catastrophe exposure by purchasing ILWs or catastrophe derivatives. Another option available to insurance and reinsurance companies is to rebalance or reduce their underwriting to lower the overall exposure to catastrophe risk.

For companies writing insurance that creates catastrophe exposure, modelling the risk of catastrophes is part of the standard business processes of underwriting and risk management; it is used also in capital allocation. Facilitating risk securitisation is not the primary goal of catastrophe modelling, even though the decision to transfer some of the risk to capital markets might be based on the modelling results. Instead, the emphasis is on total risk exposure. Modelling catastrophe risk is growing in importance at insurance and reinsurance companies, as management see the benefits it delivers. Quantification of catastrophe risk exposure is also driven by shareholders and rating agencies. Regulators are also paying more attention to catastrophe risk than ever in the past.

It would appear that the insurance industry has greater expertise in modelling catastrophe risk than the investor community. While this is generally true, there are investors who are very sophisticated in catastrophe modelling, while the insurance industry expertise is generic and not focused on the specific issues relevant to securitising insurance risk.

MODELLING AS A SOURCE OF COMPETITIVE ADVANTAGE TO INVESTORS

The primary risk of insurance-linked securities in almost all cases is, of course, the insurance risk. The risk of catastrophic events is the one most commonly transferred to investors; on the property insurance side the risk of catastrophic events fully dominates insurance securitisation. To make an informed decision, an ILS investor has to understand the risk profile of these securities. Without this understanding, it is impossible to make any intelligent decisions on individual insurance-linked securities or their portfolios. Catastrophe modelling and the risk analysis based on it are key to understanding the risk profile of these securities. (As pointed out earlier, there might be situations when an investor makes an informed decision to allocate a small portion of their assets to insurance-linked securities without developing expertise in this asset class. These situations are rare.)

Since the ability to quantify risk and determine its proper price is based

on catastrophe modelling and risk analysis, those investors better able to understand the risk analysis section of the offering circulars for catastrophe bonds have an immediate advantage over the rest of the investor community. Properly interpreting the risk analysis section requires knowledge of modelling techniques used, modelling software packages utilised, model credibility, the way exposure data is captured, and other modelling-related issues. Those who have better understanding of these issues have an advantage over those who do not. They are in a better position to quantify the uncertainty, make adjustments if necessary, and extract more useful information from the same risk analysis section of the offering circulars. This advantage is not limited to catastrophe bonds and is applicable to all types of catastrophe insurance-linked securities.

Finally, those investors who use catastrophe modelling tools themselves have an extra advantage over those who do not. They tend to have a greater degree of understanding of the assumptions underlying the models and the types of uncertainty involved. The most sophisticated of them are able to perform additional sensitivity analysis and scenario testing, to come up with a better understanding of the risk profile of the security and the price to charge for assuming this risk.

An example of the competitive advantage held by those with superior understanding of catastrophe modelling tools can be found in the analysis of California earthquake exposure. The difference in scientific views on which part of the San Andreas fault is most ripe for a major earthquake (referred to earlier in this chapter) is one of the reasons for the divergence in results among commercial catastrophe models in estimating expected losses at various exceedance levels from one part of California to another. (The divergence is true at the time of writing; models evolve, and updates and new releases are issued periodically.) Understanding the difference between models is by itself a source of competitive advantage; having an informed opinion on which model is likely to produce more precise results for a specific peril and geographical territory adds significantly to this competitive advantage. Even an informed view on the likely variability of results around the expected mean for a specific peril and geographical territory, and how it varies from model to model, is an informational advantage.

The use of models by investors is of particular importance in portfolio management. Without using real catastrophe models, all an investor can do is to make very rough estimates of the risk accumulation by peril/geography bucket and try to put limits on individual risk buckets. There is no way to properly estimate risk-adjusted return for the portfolio, or how the

addition of a position will affect the overall risk–return profile. The investors who are able to use modelling tools, both in the analysis of individual securities and in portfolio management, have an important competitive advantage, the value of which is magnified by the overall inefficiency of the insurance-linked securities market.

MODELLING AS A SOURCE OF COMPETITIVE DISADVANTAGE TO INVESTORS

The appearance of models designed specifically for investors in insurance-linked securities such as catastrophe bonds is changing the way some investors are approaching ILS investing. Some of those who never utilised catastrophe modelling tools before have now tried to use the new software to model their ILS portfolios. The models designed specifically for investors are described elsewhere, including in the chapter on portfolio management. They are much simpler to use and understand than the full-blown catastrophe models used by insurance companies and, in most cases, by modellers providing the risk analysis in structuring catastrophe bonds. They do provide ways to analyse and visualise portfolio exposure, perform “what if” analysis, and more. They appear to be simple to use.

The seeming simplicity of the tools is deceptive, however. By themselves they do not provide more than a software platform to combine individual cat bonds into one portfolio, with a semiautomatic way of calculating several risk measures. This platform is very useful to those who already understand the modelling approaches, the assumptions used in modelling, the differences between the models used for initial analysis, the degree of possible unmodelled risk, and many other factors required for using modelling tools and properly interpreting modelling results. For others, not possessing this expertise, the picture might be different. The availability of a tool that is a black box to a user can have mixed consequences. The tools themselves are not true black boxes: they are black boxes only to those who do not have the requisite expertise to use them effectively.

While most ILS investors do not use these portfolio management tools, some of those who do may be worse off than if they did not. The ability to see all securities in one portfolio and have the software spit out risk measures and other statistics can create the illusion of understanding and properly managing portfolio risk when none is present.

Modelling can be very dangerous to investors who lack the understanding of how it is performed and what the results mean. Of course, the danger is not in modelling, but in not having the level of expertise needed to understand

the modelling methods, output and implications. This problem has existed for a very long time and is unrelated to the appearance of software tools targeted specifically at the ILS investor. Improper interpretation of the risk analysis section of offering circulars by some investors has been going on for so long because of the seeming simplicity of the data presented. It creates the illusion of understanding, and that can be very dangerous. Some investors have become proficient in the lingo of catastrophe bonds and related modelling but, without realising it, have not gained the level of expertise needed to turn modelling into a useful tool. To think they understand the risk of securities when they really do not creates a dangerous situation.

The false sense of security when it comes to risk management, and the illusion of actively managing a portfolio to maximise its risk-adjusted return, can lead to catastrophic results for some investors in catastrophe risk.

One more danger to point out is that the investors focused on modelling catastrophe risk are sometimes focused on it too much, to the degree that they do not pay the necessary attention to other types of risk associated with insurance-linked securities. These other risks are important in the analysis of individual securities; it is also important to take them into account when these securities become part of an investment portfolio.

The problems mentioned above would become obvious and self-correct in investing in almost any other asset class. The level of historical returns and their volatility by itself would be a clear indicator of investor expertise, in most cases. Catastrophe ILS are tied to the risk of very rare events, and a track record of several years says little about the level of risk-adjusted returns generated.

TRENDS AND EXPECTATIONS

The importance of modelling in the analysis of insurance-linked securities is impossible to overestimate. The specific type of modelling involved in the probabilistic analysis of catastrophe events and the resulting insurance losses is unusual in the investment world and requires specialised expertise. The times when most investors made their decisions based on the rudimentary analysis of the information in the offering documents have passed. A greater level of sophistication is now required.

- Insurance and reinsurance companies seeking to transfer some of their risk to the capital markets in the form of insurance-linked securities have dramatically improved and continue to improve their risk modelling and management. They are more and more finding themselves in

the position of being able to make fully informed decisions on the ways to manage their catastrophe exposure and properly choose among such options as reinsurance, securitisation and retaining catastrophe risk.

- ❑ Superior modelling skills and the ability to better interpret results of modelling catastrophic events are a major source of competitive advantage to the investors who have this level of expertise. As the importance of modelling is becoming more widely recognised, those who lack the expertise will find it increasingly difficult to compete effectively.
- ❑ The ability to model risk is particularly valuable in assembling and managing portfolios of insurance-linked securities. This skill is even more important at the portfolio management level than in determining the right price for a particular catastrophe bond or another security whose risk is linked to catastrophic events.
- ❑ Without models, it is impossible to assess the risk-adjusted return in investing in catastrophe-linked securities. Without understanding the risk profile of a security, investors are in no position to evaluate whether they are being properly compensated for assuming the risk.
- ❑ Track record of a fund investing in insurance-linked securities can often be meaningless and even misleading. Some of the investors who have been most successful on paper have achieved higher returns by taking on disproportionate amounts of risk, often unknowingly. Without properly utilised models, we cannot analyse this type of risk. When investing in the more traditional asset classes such as equities, track record of returns is usually very informative and revealing; but it is of less importance in investing in insurance-linked securities and can be considered only in the context of the risk that has been taken. Catastrophic events are, by their very definition, very rare, and it is possible for an investor to “be lucky” for quite a long period of time even when the investment portfolio is completely mismanaged.
- ❑ An investor in catastrophe insurance-linked securities not properly using appropriate modelling tools is unable to establish an effective risk management framework around the investment process. Proper risk controls are impossible without risk modelling.
- ❑ The models are continuing to evolve and advance in their sophistication. Models for new perils and geographic regions are being developed; and, more importantly, the existing models are being improved. Better models allow better risk quantification, serving the interests of both sponsors and investors.
- ❑ Superior expertise in catastrophe modelling translates into a competitive

advantage for an investor in insurance-linked securities. It also enables better decision making for sponsors in dealing with the issues of basis risk.

- ❑ Issues of data quality, understanding model limitations, credibility of models, and biases among existing models are key components of the type of expertise that can provide a competitive advantage.
- ❑ Important as the use of modelling tools is, better understanding of the assumptions and superior interpretation of the results are of even greater significance. These two can be the most important sources of competitive advantage.

This chapter provided but an introduction to selected concepts in modelling catastrophic events in the context of analysing insurance risk securitisation. Some additional information on the topic can be found in other chapters. The issues touched on here should provide an understanding of why modelling catastrophe risk is important and why it is so difficult.

- 1 It is sometimes referred to as modified Omori law.
- 2 See J. B. Rundle *et al*, "A simulation-based approach to forecasting the next great San Francisco earthquake", *PNAS* 102(43), October 25, 2005; 15363–15367.
- 3 APEC Cooperation for Earthquake Simulation is an international project with a specific long-term goal of creating supercomputer simulation models incorporating all elements of the earthquake generation process. Similar efforts with a more narrow focus are under way at several research centres.
- 4 The data would always be provided under the conditions of confidentiality; its use is possible, if at all, by combining the data from several companies, and using it in calculations to obtain aggregate results in such a way as no confidential information is revealed.

Catastrophe Derivatives and ILWs

INDEX-LINKED CONTRACTS

Traditional insurance and reinsurance contracts are based purely on direct indemnification of the insured or reinsured for the losses suffered. Another way to transfer insurance risk, which is particularly important in its transfer to the capital markets, is to link the payments to a certain value of an index as opposed to basing it only on the reimbursement of the actual losses suffered by a specific entity. An example of such an index would be that of the level of losses suffered from a hurricane in a particular region by the whole insurance industry. Another example would be a purely parametric one based on the intensity of a specified catastrophic event without referencing actual insured losses.

The two main types of insurance-linked securities whose payout depends on an index value are insurance derivatives and industry loss warranties. (Chapter 3 describes property catastrophe bonds and Chapter 11 describes extreme mortality bonds; each of them can also be dependent on an index.) Industry loss warranties (ILWs) and catastrophe derivatives (a subset of insurance derivatives) were the first insurance-linked securities to appear. ILWs were first introduced in the 1980s and at the time they were often referred to as original loss warranties (OLWs) or original market loss warranties. The first catastrophe derivative contracts were developed in 1992 by the Chicago Board of Trade (CBOT). Both types of contract have since evolved; their markets have evolved as well. ILWs in particular are now playing an important role in the transfer of catastrophe risk from insurance to capital markets.

The use of an index as a reference offers the transparency and lack of moral hazard that are so important to investors. The ease of standardisation is also important. One of the key advantages, not yet fully realised, is the liquidity and price discovery that come with exchange-traded products such as catastrophe derivatives.

This chapter provides an overview of ILWs and catastrophe derivatives and explains the considerations used in their analysis by investors and insurers. It then describes the standard indexes used in structuring these securities and gives some specific examples. The focus is on property insurance risk transfer; insurance derivatives linked to mortality and longevity are explained in the chapters dealing with mortality and longevity risk trading, while weather derivatives are discussed in Chapter 8. Finally, the present chapter examines the trends in the market for ILWs and catastrophe derivatives and the expectations for its growth and evolution.

ROLE OF AN INDEX

Index-linked investments are common in the world of capital markets. The indexes used in insurance and reinsurance risk analysis are typically related to the level of insurance losses; these are not investable indexes and neither are their components. A derivative contract can still be structured based on such an index, but the underlying of the derivative contract is not a tradable asset.

In the transfer of insurance risk, an index is chosen in such a way that there is a direct relationship between the value of the index and the insurance losses suffered. There is, however, a difference between the two: the basis risk. This risk is not present when a standard reinsurance mechanism is utilised.

While index-linked products are used primarily for the transfer of true catastrophe risk, there is a growing trend of transferring higher-frequency (and lower-severity) risk to the capital markets. The indexes used do not necessarily have to track only catastrophic events.

CATASTROPHE DERIVATIVES DEFINED

In financial markets, a derivative is a contract between two parties the value of which is dependent on the value of another financial instrument known as an underlying asset (usually referred to simply as an underlying). A derivative may have more than one underlying. In the broader sense, the underlying does not have to be an asset or a function of an asset. Catastrophe derivatives are such contracts, with an underlying being an index reflecting the severity of catastrophic events or their impact on insurance losses.

Futures are an example of derivative instruments. Catastrophe futures are standardised exchange-traded contracts to pay or receive payments at a specified time, with the value of the payment being a function of the value

of an index. Unlike the case of traditional financial futures, physical delivery of a commodity or other asset never takes place. Options are another example of financial derivatives; they involve the right to buy (call option) or sell (put option) an underlying asset at a predetermined price (strike). In the context of catastrophe derivatives, of particular importance are call spreads, which are the combination of buying a call at a certain strike price and selling a call on the same underlying at a higher strike, with the same expiration date. The calls can be on catastrophe futures. Using a call spread limits the amount of potential payout, making the contract somewhat similar to reinsurance, where each protection layer has its own coverage limit.

Binary options provide for either a fixed payment at expiration or, depending on the value of the underlying, no payment at all. In other words, there are only two possible outcomes. They are also referred to as digital options.

There are numerous ways that catastrophe derivatives can be structured. The payout may depend on a hurricane of specific magnitude making a landfall in a certain area; on the value of total cumulative losses from hurricanes to the insurance industry over a certain period of time for a specified geographical region; or on the value of an index tracking the severity of an earthquake at several locations. The flexibility in structuring an over-the-counter (OTC) derivative allows hedgers to minimise their basis risk. At the same time, there are significant advantages to using standard instruments that can be traded on an exchange. Exchange-traded derivatives are more liquid, allow for quicker and cheaper execution, provide an effective mechanism for managing credit risk and bring price transparency to the market, all of which are essential for market growth.

Derivatives versus reinsurance

All insurance and reinsurance contracts may be seen as derivatives, albeit not recognised as such by accounting rules. Technically, they would be call spreads, which corresponds to policy limits in insurance. From the point of view of the party being paid for assuming the risk, an excess-of-loss reinsurance contract can be seen as being equivalent to selling a call with the strike at the attachment point and buying a call with the strike equal to the sum of the attachment point and the policy limit. The “underlying” in this case is the level of insurance losses.

The true derivatives such as insurance catastrophe derivatives have a better defined and stable underlying and are accounted for as financial

derivative products. Insurance accounting is not allowed for these products. This topic will be revisited later in the chapter.

INDUSTRY LOSS WARRANTIES DEFINED

The term “industry loss warranty” (ILW) has been used to describe two types of contract, one of them a derivative and the other a reinsurance contract. In its most common form, an ILW is a double-trigger reinsurance contract. Both trigger levels have to be exceeded for the contract to pay. The first is the standard indemnity trigger of the reinsured suffering an insured loss at a certain level, that is, the ultimate net loss (UNL) trigger. The second is that of industry losses or some other index level being exceeded. The index of industry losses can be, for example, the one determined by the Property Claim Services (PCS) unit of Insurance Services Office, Inc. (ISO).

An ILW in a pure derivative form is a derivative contract with the payout dependent only on the industry-based or some other trigger as opposed to the actual insurance losses of the hedger purchasing the protection. Even though labelled an ILW, it is really an OTC derivative such as the products described above.

The choice between the ILW reinsurance and derivative forms of protection has significant accounting implications for the hedger. It is typically beneficial for the hedger to choose a contract that can be accounted for as reinsurance, with all the associated advantages. This is why the vast majority of ILW transactions are done in the form of reinsurance.

The majority of ILWs have a binary payout, and the full amount is paid once the index-based trigger has been activated. (We assume that the UNL trigger condition, if present, has been met.) However, some ILW contracts have non-binary, linear payouts that depend on the level of the index above the triggering level. There seems to be general market growth in all of these categories.

MARKET SIZE

While the size of the catastrophe bond market is known, it is difficult to estimate the volume of the industry loss warranty and catastrophe derivative market. The OTC transactions are rarely disclosed, leading to a wide range of estimates of market size. The only part of the market with readily available data is that of exchange-traded catastrophe derivatives. The exchanges report the open interest on each of their products.

While its size is not very big (with no estimates exceeding US\$10 billion in limits), this market is important as a barometer of reinsurance rates and

their movements. Exchange-traded products bring price transparency to the traditionally secretive reinsurance market. The growing activity of ILW brokers is leading to increased transparency in the OTC markets as well. While not directly comparable to traditional reinsurance contracts, catastrophe derivatives and ILWs provide an important reference point in pricing reinsurance protection.

It is likely that in terms of total limits, the ILW and catastrophe derivative market is between US\$5 and US\$10 billion. This number does not include catastrophe and other insurance derivatives linked to mortality and longevity; only property and casualty insurance risks are included. The market has been growing, but the growth has not been steady. Similar to the retro market (of which some consider this market a part), its size is particularly prone to fluctuations based on the rate levels in the traditional reinsurance market. The one part of the market that we can see growing is that of exchange-traded insurance derivatives. However, exchange-traded products are currently a relatively small part of the overall marketplace.

KEY INDEXES

A number of indexes have been used in structuring insurance derivatives and ILW transactions. They include indexes tied directly to insurance losses and those tied to physical parameters of events that affect insurance losses. The overview below focuses on the indexes providing the most credible information on the level of insured industry-level property losses due to natural catastrophes.

Property Claim Services

PCS, a unit of ISO, collects, estimates and reports data on insured losses from catastrophic events in the US, Puerto Rico and the US Virgin Islands. While every single provider of catastrophe-insured loss data in the world has at times been criticised for supposed inaccuracies or delays in reporting, PCS is generally believed to be the most reliable and accurate. In the half a century since it was established, the organisation has developed sound procedures for data collection and loss estimation. It has the ability to collect, on a confidential basis, data from a very large number of insurance carriers as well as from residual market vehicles such as joint underwriting associations. Other data sources are used as well. Insurance coverage limits, coinsurance, deductible amounts and other factors are taken into account by PCS in estimating insured losses. Estimates are provided for every catastrophe – which is defined by PCS as an event that causes US\$25 million or

more in direct insured property losses and affects a significant number of policyholders and insurers. Data for both personal and commercial lines of business is included.

Loss estimates are usually reported within two weeks of the occurrence of a PCS-designated catastrophe (and PCS provides the event with a serial number). For events with likely total insured property loss in excess of US\$250 million, PCS conducts re-surveys and reports their results approximately every 60 days until it believes that the estimate reasonably reflects insured industry loss. These larger events are the ones of interest for catastrophe derivatives and ILWs. Figure 5.3 shows an example of PCS loss estimates for Hurricane Ike at various time points, in reference to the settlement prices for two of the exchange-traded catastrophe derivatives that use PCS-based triggers.

While general catastrophe loss data is available dating back to the establishment of PCS in 1949, the more detailed data by geographic territory and insurance business line is available for only the more recent years.

In Table 5.1, opposite, we can see the development of industry-insured loss estimates for the largest catastrophic events since 2001. The time between the occurrence of a catastrophic event and reporting of the final estimate could vary significantly depending on the event and complexity of the data collection and extrapolation. Of the events shown in Table 5.1, Hurricane Katrina had 10 re-survey estimates issued, with the last one almost two years after the event occurrence. However, the changes over the year preceding the reporting of the final estimate were minuscule. The 2008 Hurricane Gustav had the final estimate issued in less than five months, with that final number not changing from the first re-survey estimate.

Insured loss estimates for catastrophes that happened before those shown in Table 5.1 often lacked precision, even though they did not take longer to obtain. For the 1994 Northridge earthquake in California, the preliminary estimate increased 80% in two months, and the final estimate was five times greater than the original number. However, we have to recognise the fact that the methodologies employed by PCS have been changing; current estimation techniques are more reliable given the possibly disproportionate focus on the actual reported numbers years ago.

Catastrophe loss indexes based on PCS data are the basis for many ILW and catastrophe derivative transactions, as well as for catastrophe bonds and other insurance-linked securities. Both single-event and cumulative catastrophe loss triggers can be based on PCS indexes.

Table 5.1 Changes in PCS estimates over time for largest US catastrophic events since 2001 as reported by PCS

Year	Catastrophic event	Preliminary estimate (US\$ billions)	→ % change	First re-survey estimate (US\$ billions)	→ % change	Final estimate (US\$ billions)
2001	Wind and Thunderstorm (38-01)	0.6	193	1.7	29	2.2
2001	Tropical Storm Allison (44-01)	1.2	105	2.5	0	2.5
2001	World Trade Center – Fire-Other (48-01)	16.6	0	16.6	13	18.8
2002	Wind and Thunderstorm (61-02)	0.7	22	0.9	96	1.7
2003	Wind and Thunderstorm (88-03)	1.5	102	3.1	2	3.2
2003	Hurricane Isabel (95-03)	1.2	44	1.7	0	1.7
2004	Hurricane Charley (26-04)	6.8	0	6.8	10	7.5
2004	Hurricane Frances (28-04)	4.4	0	4.4	4	4.6
2004	Hurricane Jeanne (29-04)	3.2	6	3.4	6	3.7
2004	Hurricane Ivan (30-04)	6.0	18	7.1	0	7.1
2005	Hurricane Katrina (49-05)	34.4	11	38.1	8	41.1
2005	Hurricane Rita (51-05)	4.7	6	5.0	13	5.6
2005	Hurricane Wilma (54-05)	6.1	38	8.4	22	10.3
2008	Hurricane Gustav (58-08)	1.9	13	2.2	0	2.2
2008	Hurricane Ike (60-08)	8.1	32	10.7	17	12.5

Source: PCS

PERILS

Incorporated in 2009, PERILS AG was created to provide information on industry-insured losses for catastrophic events in Europe, similar to the way PCS provides information in the US. The plans call for ultimate expansion of catastrophe data reporting beyond Europe to other regions outside the US. The shareholders of the company are major insurance and reinsurance companies and a reinsurance intermediary, ensuring that a large segment of catastrophe loss data will be provided to PERILS. The information is provided anonymously by insurance companies and includes exposure data (expressed as sums insured) by CRESTA zone and by country, property premium data by country, and catastrophic event loss data by CRESTA zone and by country. The data is aggregated and extrapolated to the whole insurance industry based primarily on known premium volumes. Industry exposure and catastrophe loss data are examined for reasonableness and tested against information from other sources. The methodology is still evolving.

In December 2009, PERILS launched an industry loss index service for European windstorm catastrophic events. The data can be used for industry loss warranties (ILW) and broader insurance-linked securities (ILS) transactions involving the use of industry losses as a trigger. Table 5.2 provides a description of the PERILS indexes for ILS transactions.

ILW reinsurance transactions based on a PERILS catastrophe loss index have been done shortly after the introduction of the indexes. The scope and number of the indexes are expected to grow. The data collected by PERILS

Table 5.2 ILS indexes provided by PERILS AG

Index characteristic	Options or description
Covered perils	Windstorm and ensuing perils
Covered territories	Belgium, Denmark, France, Germany, Ireland, Luxembourg, Netherlands, Switzerland, UK
Line of business	Property insurance, split into residential, commercial, industrial and agricultural
Reporting schedule	First index value report at latest six weeks after the event, updated after three, six and 12 months. Subsequent reports only if deemed necessary. Reporting closed in any case after 36 months.

Source: PERILS AG

will allow the company to create customised indexes for bespoke transactions. The reporting is done in euros as opposed to US dollars.

Swiss Re and Munich Re indexes

The two largest reinsurance companies, Swiss Re and Munich Re, have been compiling industry loss estimates for catastrophic events for decades. Swiss Re's sigma, in particular, has been compiling very reliable loss estimates for catastrophe events worldwide, including manmade catastrophes. Munich Re has assembled a very large inventory of catastrophic events in its NatCatSERVICE loss database. It is similar to Swiss Re's sigma in its broad scope but does not include manmade catastrophes. Economic losses from catastrophic events are often estimated in addition to the insured losses. ILW transactions have been performed based on both Swiss Re's sigma and Munich Re's NatCatSERVICE.

It is likely that for the windstorm peril Swiss Re's and Munich Re's estimates are not going to be used for ILS transactions, since PERILS provides a credible independent alternative. Other perils, and other regions around the world usually do not have such an alternative, and it is likely that Swiss Re and Munich Re indexes will continue to be used in structuring ILW and other transactions. This practice may change in the future if PERILS implements its ambitious expansion plans.

CME hurricane index

This index has been developed specifically to facilitate catastrophe derivative trading. The index, based purely on the physical characteristics of a hurricane event, aims to provide a measure of insured losses without the use of any actual loss data such as reported industry losses. Details of the index calculation are presented in Panel 5.1. While the index has been developed for North Atlantic hurricanes, in theory the same or a similar approach can be used for cyclone events elsewhere.

Mortality and longevity indexes

A number of indexes tracking population mortality or longevity have been developed for the express purpose of structuring derivative transactions. These indexes are usually based on general population mortality as opposed to that of the insured segment of the population. They can be used for managing the risk of catastrophic mortality jumps affecting insurance companies, or the longevity risk affecting pension funds, annuity product providers and governments.

PANEL 5.1 CME HURRICANE INDEX

The CME hurricane index (CHI) was originally developed by reinsurance broker Carvill and is still usually referred to as the Carvill index. CME Group currently owns all rights to it.

The standard Saffir–Simpson hurricane scale is discrete and provides only five values (from 1 to 5) based on hurricane sustained speed. Having only five values can be seen as lacking in precision required for more accurate estimation of potential losses. In addition, the Saffir–Simpson scale does not differentiate between hurricanes of different sizes as measured by the radius of the hurricane. Hurricane size can have a significant effect on the resultant insurance losses. CHI attempts to improve on the Saffir–Simpson scale by providing a continuous (as opposed to discrete) measure of sustained wind speeds and by incorporating the hurricane size in the calculation. The following formula is used for calculating CHI

$$CHI = \left(\frac{V}{V_0} \right)^3 + \frac{3}{2} \left(\frac{R}{R_0} \right) \left(\frac{V}{V_0} \right)^2$$

V here is the maximum sustained wind speed, while R is the distance that hurricane-force winds extend from the centre of the hurricane. The denominators in the ratios are the reference values. V_0 is equal to 74 m.p.h., which is the threshold between a tropical storm and a hurricane as defined by the Saffir–Simpson scale used by the National Oceanic and Atmospheric Administration (NOAA) of the US Department of Commerce. The index is used only for hurricane-force wind speeds, that is, for V equal to or greater than 74 m.p.h. R_0 is equal to 60 miles, which is a somewhat arbitrarily chosen value intended to represent the radius of an average hurricane in the North Atlantic.

EQECAT is the current official calculation agent of the CHI for CME Group. In calculating the value of the index used for contract settlement, EQECAT utilises official data from NOAA. If some of the data is missing, which would likely involve the radius of hurricane-force winds, EQECAT is to use its best efforts to estimate the missing values. There are additional rules governing the determination of which of the public advisories (from NOAA) is to be used, what constitutes a hurricane landfall, and how multiple landfalls of the same hurricane are treated.

There is also an index tracking mortality of a specific group of individuals who have settled their life insurance policies, as opposed to the mortality of

the general population. Life-settlement mortality tracked by such an index is very different from and not to be confused with mortality of the insured segment of the population.

This chapter focuses on non-life insurance derivatives and ILWs. Mortality and longevity indexes and the insurance derivative products based on them are described in detail in the chapters dealing with securitised life insurance risk and the hedging of longevity risk.

MODELLING INDUSTRY LOSSES

Modelling losses for the whole industry is performed using the tools that are used for modelling losses for a portfolio of risks. Industry loss estimates are significantly more stable than those of underwriting portfolios of individual insurance companies. Data such as premium volume provides additional information that assists in making better predictions. In addition, using probabilistic estimates of industry losses is a natural way of comparing different modelling tools. An outlier would be quickly noticed and need to be explained. Expected annual losses for peak hazards produced by different modelling tools do not significantly diverge. The overall probability distributions, however, can differ considerably.

As an example, Table 5.3 shows estimated probabilities of insurance industry losses, as would be calculated by PCS, from a single catastrophic event exceeding a certain level that is used as trigger for catastrophe

Table 5.3 Estimated annual exceedance probabilities for industry loss damage from a single hurricane event impacting the US (all 50 states included)

Exceedance level (trigger)	Estimated exceedance probability (%)
US\$10 billion	29
US\$20 billion	15.5
US\$30 billion	9.5
US\$40 billion	6.25
US\$50 billion	4.25

Source: Navigation Advisors LLC and industry sources.

Note: No claim is made as to their accuracy of the exceedance probabilities or their applicability to a specific situation. Exceedance probabilities may vary, perhaps significantly, depending on factors such as ENSO (El Niño and La Niña). Changes to catastrophe models could lead to significant adjustments to the exceedance probabilities.

derivatives and industry loss warranties. The probabilities do not correspond directly to the results of any of the standard catastrophe models. The assumption based on significantly heightened hurricane activity and warm sea surface temperature is used instead of utilising the entire historical event catalogue. This explains the higher than usually assumed probabilities of exceedance.

THE ILW MARKET

The ILW market is very similar to the traditional reinsurance market in that it is facilitated, almost exclusively, by reinsurance brokers. The three largest reinsurance brokers, Aon Re, Guy Carpenter and Willis Re, account for almost all of the market volume. There are several small brokers that participate in the ILW market, but their share is small. Investment banks, despite their role in ILS markets in general, have limited involvement in ILWs.

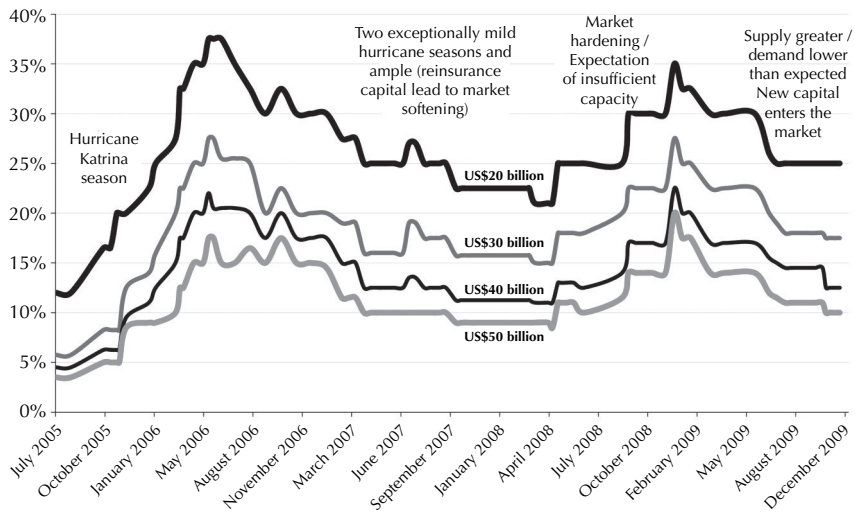
The vast majority of ILWs provide protection against standard risks of wind damage and earthquakes in the US, wind in Europe and earthquakes in Japan. All natural perils coverage for all of these territories is also common. The US territory can be split into several pieces, of which Florida has the most significant exposure to hurricane risk. In addition, second- and third-event contracts are often quoted. For these perils, in the US the standard index is PCS losses, with trigger points ranging from as low as US\$5 billion in industry losses to as high as US\$120 billion or even greater to provide protection against truly catastrophic losses.

Figure 5.1, opposite, illustrates indicative pricing for 12-month ILWs covering the wind and flood risk in all of the US. The prices, expressed as a percentage of the limit, are shown for first-event contracts at four trigger levels: US\$20 billion, US\$30 billion, US\$40 billion and US\$50 billion. The trigger levels are chosen to correspond to those used later in the chapter in the illustration of price levels for the IFEX contracts covering substantially the same catastrophe events.

The prices can be seen to fluctuate dramatically depending on the market conditions. The highest levels were achieved following the Katrina–Rita–Wilma hurricane season of 2005. Another spike followed the 2008 hurricane losses combined with the capital depletion due to the financial crisis. The expectations of even higher rates immediately before the hurricane season of 2009, however, did not materialise.

Structuring an ILW

Industry loss warranties have become largely standardised in terms of their typical provisions and legal documentation. A common ILW agreement will

Figure 5.1 US Windstorm ILW indicative pricing

Note: All contracts have the duration of 12 months. Figures based on average of broker indicative pricing when available.

be structured to provide protection in case of catastrophic losses due to a natural catastrophe such as a hurricane or an earthquake.

The first step will be deciding on the appropriate index, which in the US can be a PCS index. Once the index is chosen, the attachment point has to be determined, as well as the protection limit. As the value of losses from a catastrophic event is not immediately known and an organisation such as PCS will need time to provide a reliable estimate, a reporting period needs to be specified to allow for loss development. This period can be, for example, 24 months from the date of the loss or 18 months from end of the risk term. The contract risk term is generally 12 months or shorter. Some ILWs provide protection only during the hurricane season. For earthquake protection, the 12-month term is standard. Multi-year contracts are rare.

As an example of the legal language in a contract providing protection against catastrophic losses due to an earthquake, the contract might “indemnify the Reinsured for all losses, arising from earthquake and fire following such earthquake, in respect of all policies and/or contracts of insurance and/or reinsurance, including workers’ compensation business written or assumed by the Reinsured, occurring within the territorial scope hereon. This Reinsurance is to pay in the event of an Insured Market Loss for property business arising out of the same event being equal to or greater than

US\$20 billion (a 'Qualifying Event'). For purposes of determining the Insured Market Loss, the parties hereto shall rely on the figures published by the Property Claim Services (PCS) unit of the Insurance Services Office." The US\$20 billion is specified as an example of the trigger level.

The limits can be specified in the manner typical of an excess-of-loss reinsurance contract, with the possible contract language stipulating that the reinsured will be paid up to a certain US dollar amount for "ultimate net loss each and every loss and/or series thereof arising out of a Qualifying Event in excess of" an agreed-upon "ultimate net loss each and every loss and/or series thereof arising out of a Qualifying Event". A reinstatement provision usually would not be included, but there are other ways to assure continuing protection after a loss event, including purchasing second- or multiple-event coverage, which can also be in the form of an ILW.

While the reinsurance agreement requires that both conditions be satisfied – that is, only actual losses be reimbursed and only when the industry losses exceed a predetermined threshold – the agreements tend to be structured so that only the latter condition determines the payout. The attachment point for the UNL is generally chosen at a very low level, ensuring that exceeding the industry loss trigger level will happen only if the reinsured suffers significant losses. There is, however, a chance of the contract being triggered but the covered UNL being below the full reinsurance limit.

Arguably the most important element of an ILW contract is the price paid for the protection provided. The price would typically be expressed as rate on line (RoL), that is, the ratio of the protection cost (premium) to the protection limit provided. The payment is often made upfront by the buyer of the protection.

An important issue in structuring an ILW is management of credit risk. This topic is covered later in the chapter. Collateralisation, either full or partial, might be required to assure payment. The need for collateralisation is more important when the protection is provided by investors as opposed to a rated reinsurance company.

ISDA US WIND SWAP CONFIRMATION TEMPLATE

In 2009, the International Swaps and Derivatives Association (ISDA) published a swap confirmation template to facilitate and standardise the documentation of natural-catastrophe swaps referencing US wind events. Prior to that, several templates existed in the marketplace. The ISDA template is based on the one originally developed by Swiss Re. The template

uses PCS estimates for insurance industry loss data for catastrophic wind events affecting the US. The covered territory is defined as all of the US, including the District of Columbia, Puerto Rico and US Virgin Islands. The option of choosing a subset of this territory also exists. It allows the choice of three types of covered event: USA Wind Event 1, USA Wind Event 2 and USA Wind Event 3. The first type is the broadest and includes all wind events that would be included in the PCS Loss Report. The second specifically excludes named tropical storms, typhoons and hurricanes, while the third includes only named tropical storms, typhoons and hurricanes. As in all of the swap confirmations used in the past for US wind, flood following covered perils is included in the damage calculation. The template clarifies the treatment of workers' compensation losses, and whether loss-adjustment expenses related to such losses are included. It allows for both binary and non-binary (linear) payments in the event of a covered loss.

The ISDA template specifically states that the transaction is not a contract of insurance and that there is no insurable loss requirement. The structure is that of a pure financial derivative without any insurance component.

While the template brings legal documentation standardisation to these OTC transactions, it allows a significant degree of customisation to minimise the basis risk of the hedging party; this degree of customisation is not possible when using only exchange-traded instruments.

IFEX CATASTROPHE DERIVATIVES

Of the exchange-traded catastrophe derivatives, IFEX event-linked futures (ELF) are one of the two most common, the other being CME catastrophe derivatives. IFEX is the Insurance Futures Exchange, which developed (together with Deutsche Bank) event-linked futures. IFEX event-linked futures are traded on the Chicago Climate Futures Exchange (CCFE), a relatively new exchange focused on environmental financial instruments. CCFE is owned by Climate Exchange PLC, a UK publicly traded company. The founder of CCFE, Richard L. Sandor, played a key role in the introduction of the first catastrophe derivative products in the early 1990s. Even though the products were well designed, at the time the insurance industry was not ready for such a radical innovation as trading insurance risk. In addition to the need for education, the industry then did not have proper tools to quantify catastrophe risk or to estimate the level of basis risk created by the use of index-linked products as opposed to traditional reinsurance.

The CCFE IFEX contracts have been designed to replicate, as far as possible, the better-known and accepted ILW contracts. The two primary

differences between a traditional ILW and the corresponding IFEX contract are, first, that IFEX event-linked futures are financial derivatives and not reinsurance, and, second, that IFEX contracts provide an effective way to minimise if not eliminate the counterparty credit risk present in many ILW transactions. The terms “IFEX contract” and “ELF contract” are often used interchangeably.

IFEX contract specifications

There are currently the following types of PCS-based contract for the wind peril, which differ by the territory they cover:

- ☐ US Wind (all 50 states and including Alaska and Hawaii, Puerto Rico, US Virgin Islands and Washington, DC);
- ☐ Florida Wind (Florida only);
- ☐ US Gulf Coast Wind (Alabama, Mississippi, Louisiana and Texas);
- ☐ US Eastern Seaboard Wind (seaboard states from Georgia to Maine); and
- ☐ US North East (seaboard states from Virginia to Maine).

Most of these contracts have not been traded and were introduced only recently. The activity has been concentrated on the US Wind contracts, and, to a lesser extent, on the Florida Wind contracts. Florida Wind is the main component of the US Wind contracts.

Key specifications of US Wind IFEX event-linked futures are presented in Table 5.4. Each IFEX contract has the notional value of US\$10 thousand. The event claim index varies from 0 to 100; multiplying the value of the index by US\$100 (as per Table 5.4) can produce the maximum value of US\$10,000.

Prices for IFEX contracts have at times exhibited idiosyncratic behaviour, in part due to the insufficient liquidity that is common to all new products. Figure 5.2 shows settlement price changes over time for the 2009 first-event US windstorm at four different trigger levels.

Settlement prices are established by the exchange twice a day. Since the trading volume for this new product is light, the settlement price is not necessarily equal to the price at which the latest transaction has been performed. The bid–offer spread is rather wide for some contracts, while for some others there might not be any quotes at all at a particular time. The exchange often uses a significant degree of judgement in determining settlement prices to assure general reasonableness and consistency across trigger levels.

Figure 5.3 shows prices for IFEX contracts that were exposed to losses

Table 5.4 US Tropical Wind IFEX event-linked futures (ELF) main contract specifications

Contract parameters	Specifications
Contract Size	US\$100 multiplied by Event Claim Index.
Quotation Currency	US\$
Minimum Tick Increment	0.05 Event Claim Index point per contract = US\$5 per contract.
Contract Listing Cycle	Minimum of two annual December contract series. Each contract has its risk period of January 1–December 31 of the contract year.
Industry Loss Reporting Service	PCS.
Covered Event	<p>A “Covered Event” will be deemed to have occurred with respect to any listed Loss Trigger Level when the Exchange confirms that on or before the Contract expiration for an Event Claim:</p> <ul style="list-style-type: none"> (i) a final PCS Report has been issued that reports an Industry Loss Amount resulting from an Eligible Event in an amount equal to or in excess of the applicable Loss Trigger Level for such Event Claim; or (ii) as of the Contract expiration a final PCS Report has not been issued with respect to an Eligible Event, the most recent interim PCS Report that has been issued indicates an Industry Loss Amount resulting from such Eligible Event in an amount equal to or in excess of the applicable Loss Trigger Level for an Event Claim.
Loss Trigger Level	<p>Within any listed Contract, the Exchange may offer the following Loss Trigger Products covering January 1 through December 31 of the applicable contract year:</p> <p>US\$10 billion; US\$15 billion; US\$20 billion; US\$25 billion; US\$30 billion; US\$40 billion; US\$50 billion; US\$60 billion; US\$75 billion; and US\$100 billion.</p>
Event Claim	At least one Event Claim will exist for each Loss Trigger Product. The Exchange may list additional Event Claims for any Loss Trigger Products.

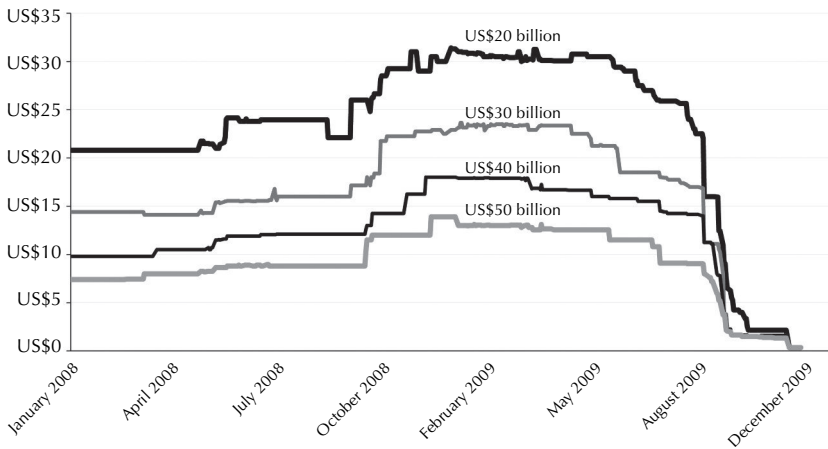
Table 5.4 *Continued*

Contract parameters	Specifications
Eligible Event	A “US Wind Event” occurring in or affecting the 50 states of the United States, Washington, DC, Puerto Rico or the US Virgin Islands (the United States Covered Territory) that has a Date of Loss falling within the Contract Risk Period for the applicable contract.
First Trading Day	An annual December contract is listed on the first business day after November 30.
Last Trading Day	The scheduled last trading day for any listed contract is the last trading day of the 18th calendar month following the end of the Contract Risk Period for the listed contract. The Exchange may declare a Last Trading Day for a listed contract earlier than the scheduled Last Trading day under certain circumstances.
Cash Settlement	Positions at each Loss Trigger Level of each Event Claim are cash-settled at Contract Expiration at an index value of either one hundred (100.00) if a Covered Event has been associated therewith, or zero (0.00) if no Covered Event has been associated therewith.
Price Limits	No daily price limits.

Sources: Chicago Climate Futures Exchange and Insurance Futures Exchange Ltd.
Contract specifications and related rules are subject to revision. Complete specifications are provided only in the CCFE Rulebook.

from Hurricane Ike in 2008. The first-event US wind contract at the US\$10 billion trigger level ultimately settled at 100 (full payment) when PCS came with its final loss estimate in October 2009. The price movements along the lifetime of the contract are instructive – in particular, changes starting shortly before Ike made a landfall and ending when the consensus was developed that losses had exceeded US\$10 billion. Each insured loss estimate issued by PCS can be seen as it is reflected in the contract price. It was almost immediately clear that insured losses from the event would not reach US\$20 billion, and the price for the first-event US\$20 billion level quickly drops as the hurricane season runs out of steam.

Figure 5.2 US Windstorm IFEX (ELF) settlement prices for first-event 2009 contracts



Note: The prices are quoted at 0.05 Event Claim Index point per contracts, that is, US\$5.00 per contract. (Contract size is equal to US\$100 times the Event Claim Index.) Settlement prices are set by the exchange and may differ from the prices at which the latest trades have been conducted. Only four trigger levels are shown. No reinstatement provision is included in IFEX contracts.


Margin requirements

As exchange-traded futures, IFEX contracts are subject to maintenance requirements. The concept of margin is unfamiliar to many insurance professionals, even though similar tools are sometimes used in the traditional reinsurance contracts. Margining makes the cashflows of both the buyer and the seller of protection different from what they would be for a reinsurance contract. The two obvious implications concern contract pricing and liquidity considerations.

There are two types of margin that have to be posted: maintenance margin and variation margin. Maintenance margin is posted by both buyers and sellers, and is intended to ensure that the parties fulfil their financial obligations under the contracts. Variation margin is simply the payment reflecting a change in the contract price: if the price increases, the seller pays to the buyer the amount equal to the price change (or, rather, the seller's account is debited and the buyer's account is credited with this amount); if the price decreases, the buyer pays the corresponding amount to the seller (the buyer's account is debited and the seller's account is credited with this amount).

The current maintenance margin requirements imposed by the CCFE are

Figure 5.3 Pricing and eventual settlement of the 2008 US Wind IFEX contract with the trigger levels of US\$10 billion and US\$20 billion



shown in Table 5.5. The term “initial margin” is not used in reference to IFEX contracts; the initial margin would always equal the maintenance margin.

It may appear counterintuitive that the seller’s margin increases if a moderate hurricane threat is declared but then decreases if the threat level is upgraded to severe. The decrease in maintenance margin when the threat level is upgraded does not imply positive cashflows to the seller. While the maintenance margin may decrease, in all likelihood this decrease is more than offset by the variation margin due to the jump in the contract price. If the threat passes, this variation margin flow is reversed, with payments made to the seller.

It is important to note that incorrect maintenance margin rules for these contracts have been circulated and can be found in some presentations posted on the Internet. Caution should be used in determining margin requirements at various time periods; the exchange is the best source of information on margin requirements.

Maintenance margin values shown in Table 5.5 are the ones established by the Chicago Climate Futures Exchange. A broker may establish higher margin requirements for some of its clients depending on the assessment of their credit risk profile. These requirements may change during the lifetime

Table 5.5 Maintenance margin requirements per contract for US Tropical Wind IFEX event-linked futures

Time	Event number	Posted by seller (US\$)	Posted by buyer (US\$)
Before June 1 of the contract year and after November 30 (unless hurricane has made landfall and the Clearing Corporation establishes other margin requirements)	Event 1 loss trigger level	US\$800 or 8% of notional amount	US\$200 or 2% of notional amount
	Event 2, 3, 4 loss trigger level	US\$400 or 4% of notional amount	US\$100 or 1% of notional amount
Between June 1 and November 30 of the contract year (absent hurricane threat or landfall)	Event 1 loss trigger level	US\$2,400 or 24% of notional amount	US\$600 or 6% of notional amount
	Event 2,3,4 loss trigger levels	US\$800 or 8% of notional amount	US\$200 or 2% of notional amount
Moderate hurricane threat is declared	Event 1 loss trigger level	US\$3,000 or 30% of notional amount	US\$3,000 or 30% of notional amount
	Event 2,3,4 loss trigger levels	US\$800 or 8% of notional amount	US\$200 or 2% of notional amount
Severe hurricane threat is declared	Event 1 loss trigger level	US\$2,000 or 30% of notional amount	US\$8,000 or 80% of notional amount
	Event 2,3,4 loss trigger levels	US\$800 or 8% of notional amount	US\$200 or 2% of notional amount
Hurricane makes a landfall	All trigger levels	Margins remain at pre-landfall level (threat level) with later adjustments made based on PCS loss estimates and at the Clearing Corporation discretion	

Source: Chicago Climate Futures Exchange.

The exchange may change the margin requirements in the future.

of a contract even if they initially equal those established by the exchange. The risk of this change should be taken into account in pricing by counterparties with potential credit problems.

While the exchange allows maintenance margin to be posted by depositing cash, high-grade securities or letters of credit, a broker might have different requirements for its clients. Of course, one of the ways of satisfying maintenance margin requirements, in whole or in part, can be through the accumulation of variation margin.

The buyer's cumulative cash outflows will never exceed the initial price of the contract. The seller's maximum cash outflow is the difference between the contract's notional amount and the initial price.

Block trades

By volume, most of the event-linked futures trades conducted on the exchange have been block trades, that is, privately negotiated transactions between two parties. The minimum size for a block trade of IFEX contracts on the CCFE is very small, at only 25 contracts. There are specific rules governing block trades; for example, a trade has to be reported to the exchange within 15 minutes of finishing the negotiation, or, if this happens to be outside the regular trading hours, within 15 minutes of the start of the next trading session. Any block trade has to be first approved by the exchange. At that moment, the longs and the shorts become subject to the margin rules of the exchange.

CME products described below have been traded in blocks as opposed to on the screen. It is possible that live quotes and on-the-screen trading will develop for CME hurricane derivatives as well at some point in the future. IFEX has the greatest on-the-screen liquidity and trading volume. CME has the same fifteen-minute reporting requirement for block trades for hurricane derivatives, which is different from the more standard five-minute rule for most products.

The fact that most of the volume comes from block trades can be seen as a negative, since block trades, unlike on-the-screen trades, do not significantly contribute to the liquidity badly needed in this market. On the other hand, there are also some positives: notably, the exchange serves as a clearing mechanism, offering an efficient way to conduct the transaction, with the credit risk being almost negligible.

CME HURRICANE DERIVATIVES

CME Group launched its hurricane derivatives products in 2006 using, instead of PCS losses, the CME Hurricane Index (at the time called Carvill Hurricane Index) described earlier. CHI provides a better measure of the destructive potential of a hurricane event than the standard hurricane scales, while allowing for quick reporting of the index values (which is not possible

when a PCS-type index is used). The choice of the index allows CME products to better address the basis risk issues of hedgers. At the same time, its more customised nature does little to attract investors – with the exception of the dedicated ILS funds and reinsurance companies, who understand the risk better and are also willing to live with the lack of liquidity. The latter statement is not meant to be criticism of the product but is a testimony to the difficulty of introducing such products to the capital markets. Direct comparisons between CME hurricane products and those traded on the CCFE are difficult. As the volume of transactions grows – as all the participants are hoping – greater liquidity can find its way to this segment of the marketplace as well.

CME hurricane derivatives have generally been negotiated off the exchange and then cleared through CME. This is very much a reinsurance broker market, with Tradition Re (part of TSF) playing a major role in arranging transactions.

Contract types

The three types of contract offered each cover one of the following:

- ☐ named storms;
- ☐ seasonal accumulated value; and
- ☐ seasonal maximum value.

For each of the three, standard futures contracts are offered based on the CHI index. In addition, options on the futures are offered as well, covering the above-mentioned three types of futures. The latest product introduced by CME is binary options on the futures. These are offered on the seasonal accumulated value and seasonal maximum value, all also based on the CHI index. The binary products are intended to more closely replicate industry loss warranties (ILWs), whose payout is almost always binary. In summary, the three products are:

- ☐ standard hurricane index futures;
- ☐ vanilla options on the futures; and
- ☐ binary options on the futures.

Seasonal maximum contracts can also be taken on the basis of the second event. Binary options of the second-event futures have been cleared through the exchange. For named storms, CME would typically want to have three contracts issued at any time, so that the fourth would be added after the occurrence of the first hurricane, and so on.

The contract size is US\$1,000 times the value of the respective CME Hurricane Index. The tick size is 0.1 CHI Index Point while the tick value (0.1 CHI Index Point) is equal to US\$100. For the binary contracts, the contract size is US\$10,000 times the value of the respective CME Hurricane Index, but the tick value is 0.01 CHI Index Point, or US\$1. Binary options pay if they are in or at the money (respective CHI value is equal to or greater than the strike). The options are all American-style and can be exercised any time up to and including the last trading day (LTD). Trading of the futures and options terminates on the first business day of the exchange following at least two calendar days after the end of the referenced calendar year. In the event of a named storm, corresponding contracts terminate on the first exchange business day following at least two calendar days after the last forecast/advisory issued by the National Hurricane Center for this named storm. Numerous additional rules apply.

Detailed description of CME hurricane contracts is not provided here; instead, the CCFE event-linked futures described above provide an illustration of exchange-traded property catastrophe derivative products.

Geographical regions

The geographical regions for CME hurricane contracts are the following:

- ☐ Cat-In-A-Box – Galveston–Mobile area bounded by 95°30'0"W on the west, 87°30'0"W on the east, 27°30'0"N on the south, and the corresponding segment of the US coastline on the north;
- ☐ Eastern US – Brownsville, TX to Eastport, ME;
- ☐ Florida – AL/FL border to Fernandina Beach, FL;
- ☐ Gulf Coast – Brownsville, TX to AL/FL border;
- ☐ Gulf Coast and Florida – Brownsville, TX to Fernandina Beach, FL;
- ☐ Northern Atlantic Coast – NC/VA border to Eastport, ME;
- ☐ Southern Atlantic Coast – Fernandina Beach, FL to NC/VA border; and
- ☐ Miami – Card Sound Bridge, FL to Jupiter Inlet, FL.

Not all territories might be available for all types of contract. The fact of the contracts being offered does not mean that all or most of them have ever been sold or bought. The Miami territory is the latest added to the list and represents the greatest risk exposure over a small geographical area in terms of potential insurance losses.

The Cat-In-A-Box region stands apart from the others, in that it is situated offshore. The losses from a hurricane hitting this region would result mostly

from the damage to oil rigs in that area and the cost of forced evacuations and shutdowns. Hurricane derivatives for this region are of interest to energy traders, who might consider combining them with positions in such products as natural gas futures. It is not necessary for a hurricane landfall to occur in order for a Cat-In-A-Box contract to be triggered.

The main CME hurricane contracts have a very narrow focus on the areas most prone to suffering extreme damage from hurricanes; these areas also often suffer from the lack of reinsurance capacity, at least at the cost considered reasonable by the buyers of protection. The focus may be narrow but is chosen to address perceived demand. The growing pains of the products have to do more with the difficulties of introducing any new product that requires the education of market participants, the need for liquidity to attract more investors and the need for analytical expertise not possessed by most investors. Additionally, as is common to all derivative products, unfavourable accounting treatment of this type of hedge factors into the growth trend.

Other considerations involving CME hurricane products

It goes without saying that only call options are available. Due to the low number of transactions on the exchange, settlement prices appear to be based primarily on the mark-to-model approach. Settlements prices are established by the exchange and are not necessarily the prices at which the last transactions were done.

This product, in particular for the Cat-In-A-Box region, is of interest to natural-gas traders. Insurance companies have not flooded to this market as was initially hoped by the exchange when the product was first introduced.

Even though the index used is not based on actual insurance losses, unlike PCS-type indexes, its use might serve the purpose of minimising basis risk if expert modelling is performed by the hedger. Attachment points of the protection coverage can often be placed lower, and the use of the right combination of the contracts covering individual territories can provide efficient protection. This is the potential of the product; though this potential has not yet materialised with transaction volume still low.

Another important difference between CME derivatives based on the CME Hurricane Index and products such as event-linked futures listed on CCFE is that CME products can be settled much faster. There is no need to wait for insurance industry loss estimates to be issued (which can take a very long time, as described above): instead, only calculations based on a known formula are needed. This expediency reduces the uncertainty factor

while also preventing the margin from being unnecessarily tied up for prolonged time periods.

The CME catastrophe products can be useful in live cat trading. They can, assuming the liquidity is present, provide for a means of last-minute hedging as well as opportunistic investing when a hurricane is approaching.

EUREX HURRICANE FUTURES

In 2009, Eurex, the largest European derivatives exchange, entered the catastrophe derivatives market by introducing hurricane futures for US hurricane risk. Binary contracts based on PCS-reported estimates of insurance industry losses were introduced for the following three regions:

- ❑ US – all 50 states, Washington, DC, Puerto Rico and US Virgin Islands;
- ❑ Florida – all the State of Florida; and
- ❑ Gulf – States of Alabama, Louisiana, Mississippi and Texas.

The contracts covering all of the US were offered at five strikes, US\$10 billion, US\$20 billion, US\$30 billion, US\$40 billion and US\$50 billion; for Florida contracts the strike levels are US\$30 billion, US\$40 billion and US\$50 billion; and for the Gulf contracts the strike levels are US\$10 billion and US\$20 billion. Only first-event contracts were introduced. If this market takes off, it is likely that the product offering will be expanded.

It is difficult to see the differences between the IFEX contracts listed on the CCFE and those introduced by Eurex. The CCFE has a broader product offering in terms of territories, strike levels and second- and subsequent-event futures, but for the same products specifications read almost exactly the same. (Some have assumed that the products are absolutely identical; in fact, small differences do exist and should be considered by a trader in these derivatives.) It is not uncommon for largely the same financial products to be traded on more than one exchange. In this sense, the introduction of hurricane futures by Eurex is not unexpected. Unfortunately, at this point the volume of trading is so light that it is difficult for even one exchange to generate profits off these products. However, even if on-the-screen trading does not develop on Eurex, the exchange can act as a clearing mechanism for trades negotiated off the exchange but with the parties wanting to take advantage of the extremely low credit risk of exchange-cleared transactions.

MORE UNUSUAL PRODUCTS

Other derivative products of a similar nature appear and disappear. An example of a such a product is Hurricane Risk Landfall Option or HuRLO, intended to bypass the insurance market and offer protection directly to businesses and homeowners in hurricane-prone zones. (The developers of the product state explicitly that the product is not intended to replace homeowners, business interruption or flood insurance.)

Developed by Weather Risk Solutions (WRS), the product was introduced as a commodity option traded on an electronic trading platform operated by WRS through the CME Alternative Marketplace's exempt board of trade. HuRLOs are similar to European-style call options, with their payout dependent on whether and where among the covered territories a hurricane makes a landfall. There is a HuRLO associated with 78 coastal counties or regions with high exposure to North Atlantic hurricanes. In addition, there is a 79th HuRLO corresponding to the case when no hurricane strikes any of the 78 territories in a given year. Series 1 of the HuRLOs covers the occurrence of the first hurricane landfall in one of the 78 territories. Series 2 covers the second hurricane landfall in the same year. The total number of HuRLOs for both Series 1 and Series 2 is then 158. An unlimited number of HuRLOs can be purchased for each outcome.

The unusual feature of the product is that it functions as a mutualised risk pool, as opposed to having a buyer and seller for every transaction. There is no need to find a counterparty to be on the other end of a transaction. After being initially seeded in return for an equal number of each of the 79 HuRLOs in a series, prices for each of the HuRLOs are set based on historical probabilities of hurricane landfalls for individual HuRLOs. As buyers purchase the HuRLOs, prices adjust based on market demand as determined by the previous transactions, subject to some restrictions concerning risk concentration and pricing stability. (Prices are established formulaically on the basis of an adaptive control algorithm that takes into account market probabilities for each of the 79 outcomes in a series.) With the exception of administrative fees and a certain percentage paid to the seed capital provider, all the premiums collected are aggregated in one fund. The total pool is then paid to buyers of one of the HuRLOs. For example, if a hurricane strikes Miami, owners of the Miami-Dade County HuRLOs receive the payout that is split among them based on the number of HuRLOs sold for that region. If no hurricane strikes any of the HuRLO regions during the calendar year, the total payout goes to the 79th ("no landfall") HuRLO in Series 1.

Option exercise is automatic. There also exists a platform for secondary

market trading. While purchasing any of the first 78 HuRLOs is generally a hedging tool, buying the 79th HuRLO (“no landfall”) is a speculative investment.

Products such as HuRLO tend to be introduced, then to disappear, and sometimes to be relaunched. They find it very difficult to get traction for a variety of reasons, primarily the difficulty in marketing them when easier-to-understand insurance solutions are available. In case of HuRLO, there are also serious concerns about the basis risk for the hedgers: the hedge effectiveness in most cases is rather low.

COMMENTS ON PRICING

Pricing ILWs and exchange-traded catastrophe derivatives is based on modelling index values as described above. In particular for exchange-traded derivatives, a proper cashflow model should be built to account for changes in margin over the life-time of the contract. Since the cashflows heavily depend on external events (such as hurricane landfalls, threat levels that change margin requirements and so forth), many scenarios should be modelled. Such a probabilistic cashflow model would most adequately address the pricing requirements for both the seller and the buyer of protection. Specific weights can be assigned to individual scenarios based on judgement in addition to the modelling output. (It is not possible to simply use an existing catastrophe model for this purpose since a number of parameters are controlled by the exchange.)

Properly taking into account cashflows due to margining of the contracts is not always done. Instead, those with background in reinsurance but not capital markets sometimes focus on the RoL and use it as the primary or sole determinant of prices. Needless to say, this approach is incorrect, although it can serve as the first approximation.

Hedge effectiveness is another point to consider in pricing these instruments from the point of view of the hedger. If the hedge effectiveness is not sufficiently high, the protection is not worth as much to the hedger. In addition to hedge effectiveness, such considerations as accounting treatment of the transactions, and the effect it has on risk-based capital, economic capital or the capital required to maintain a certain rating, all affect how much the hedger would be willing to pay for the protection. Since from the point of view of the protection seller these considerations are largely moot, in theory they should not have a significant effect on the pricing levels for the securities. However, this market is not efficient by any analysis, and these considerations do play an important role in setting price levels and in the supply–demand dynamics.

CREDIT RISK

Mitigation of counterparty risk has grown in importance after the Lehman default and other events of 2008. Collateralisation has become more important, and the quality of collateral more closely scrutinised.

There seems to be limited uniformity in how the collateral issue is handled in ILW transactions. Since many ILW protection sellers are reinsurance companies, the rating might be sufficient to alleviate credit risk concerns. In these cases, no collateral might be required. In some cases, partial collateral might suffice, or posting collateral might be required only during the hurricane season. In most cases, the restrictions on the types of assets in collateral accounts are less stringent than those found in such types of insurance-linked securities as catastrophe bonds. It appears that some of the protection buyers are much less demanding than others in issues of collateralisation.

Pure investors (as opposed to reinsurance companies) find themselves at a disadvantage in these transactions since they rarely have a credit rating and typically have to post full collateral from day one of the contract period. For this reason, few pure investors have recently been providing ILW protection. (We do not consider dedicated ILS funds that are active in reinsurance to be pure investors in this sense.) The playing field is perceived to be uneven, but this situation is likely to change.

Exchange-cleared products provide the protection against credit risk that eliminates the need for collateralisation. For example, the counterparty for all transactions involving event-linked futures on the Chicago Climate Futures Exchange is the Clearing Corporation, now part of Intercontinental Exchange (ICE). ICE operates regulated global futures exchanges and OTC markets for numerous products. ICE US Trust, LLC (ICE Trust) is a member of the Federal Reserve System and a clearing house and central counterparty for many types of transaction. The margining system serves the purpose of minimising the risk of default by the Clearing Corporation. The credit risk of exchange-cleared transactions is remote.

Additional discussion of credit risk issues in insurance-linked securities can be found in Chapter 7.

BASIS RISK

The issue of basis risk is often raised in connection with index-linked products such as ILWs and exchange-traded derivatives. There is always a chance of significant losses to the hedger if the index-linked product does not provide payment as intended. For PCS-type indexes, this risk is a

function of how different the underwriting portfolio of the hedger is from that of the insurance industry as a whole. Strike level (attachment point) and the types of ILW or exchange-traded products have to be set based on careful modelling to increase hedge effectiveness.

Exchange-traded derivatives and ILWs done in the derivative form can also lead to a situation where the hedger is paid even though it has not suffered significant losses. This can happen even for ILWs in the reinsurance form, since the insured loss trigger (as opposed to the index-based one) is usually set at a very low level.

UNL reinsurance coverage does provide protection with minimal basis risk, and is usually the first choice of protection for buyers. The reasons for entering into ILW or catastrophe derivative transactions have to do with other considerations such as lack of reinsurance (and in particular retrocessionary) capacity at affordable prices, which override the basis risk concerns.

THE USE OF TRANSFORMERS

In some cases, the insurance or reinsurance company seeking to hedge its risk by purchasing a catastrophe derivative would prefer to have the transaction accounted for as reinsurance. There are a number of benefits in the reinsurance accounting treatment that are unattainable in derivative transactions. To avoid this difficulty, a transformer structure is often utilised. It does exactly what its name implies: it acts as a transformer between reinsurance and investment. A transformer could be a separate reinsurance company (in practical settings a segregated account or “cell”) that provides fully collateralised reinsurance protection. The collateral comes from investors who purchase non-voting preferred shares in the company. (Other structures can be used as well.) Sometimes a reinsurance company will decide to assume the role of a transformer by using its general account. It can then hedge the risk by entering into a derivative transaction and retain the basis risk.

Using a transformer adds to the cost of the transaction, but for many insurance and reinsurance companies it is still the most efficient way to obtain protection. There have been transformers set up for the express purpose of allowing reinsurance accounting for exchange-cleared derivative transactions.

There could also be reasons why the protection seller would want to structure the transaction in the reinsurance form, but this happens very rarely.

Using a transformer for an exchange-traded product may be seen as defeating the purpose of the exchange-traded catastrophe derivative market

since it does little to contribute to the liquidity of the derivatives. “Transformed” derivatives are not traded. It may be that the accounting rules will be changed, eliminating the need for a transformer, but at this point such a change seems unlikely.

INVESTOR UNIVERSE

ILWs and catastrophe derivatives provide investors with one more tool for assembling and optimising an investment portfolio. However, these instruments are less understood by the investor community than, for example, catastrophe bonds. There are fewer shortcuts in the investor analysis, as these securities do not have a credit rating. The full analysis has to be performed. At the same time, some of the instruments are easier to analyse than catastrophe bonds, since probabilistic modelling looks at industry losses instead of losses to a specific underwriting portfolio, so several layers of uncertainty are removed. That said, proper modelling of these securities in the portfolio context presents largely the same challenges as analysing any insurance-linked security.

The nature of these instruments tends to limit the investor universe to specialists who are better able to analyse ILWs and catastrophe derivatives and who possess the necessary expertise. Many of the sellers of protection in this market are reinsurance companies that take advantage of their catastrophe-modelling capabilities. Dedicated ILS funds also play an active part; they sometimes understand the risk better than the reinsurance companies, even though their analytical resources are not as great. Both reinsurance companies and the dedicated ILS funds can also be on the other side of the transaction by purchasing protection to manage their portfolios in the most efficient manner.

There are also investors who do not generally invest in insurance risk but find insurance derivatives an effective way to gain exposure to this asset class and the diversification that comes with it.

In addition, some protection buyers could come from outside the insurance industry and be energy traders or investors with significant real estate holdings in hurricane-prone areas. Commodity traders might want to consider catastrophe derivatives as part of their hedging programme.

MORTALITY AND LONGEVITY DERIVATIVES

As mentioned above, there exist a number of indexes tracking general population mortality and longevity as well as those that reflect only specific populations (such as the insured who have settled their insurance policies).

Derivative products based on these indexes serve the purpose of transferring or investing in the risk of mortality spikes or longevity being higher than expected.

INVESTOR AND HEDGER PERSPECTIVES

The lines between a capital markets participant investing in insurance risk (that is, providing additional capacity to the reinsurance markets) and a hedger purchasing protection against catastrophic risk can be blurred in the case of ILWs and catastrophe derivatives. First, many or even most protection sellers are reinsurance companies; while others are often dedicated ILS funds with their own reinsurance operations. Second, the purchasers of protection include not only insurance and reinsurance companies but also dedicated ILS funds that actively manage their portfolios.

While ILWs and catastrophe derivatives introduce basis risk for the hedger, they can be cheaper than traditional reinsurance solutions that avoid this risk. The collateralisation reduces credit risk; in the case of exchange-traded derivatives, credit risk is almost completely eliminated.

Transactions based on an index are cheaper to execute. They also bring to the market a degree of standardisation that tends to put downward pressure on prices.

Index-based transactions are usually easier to model, which has the potential to attract a broader universe of investors to this market. Again, more investor capital is in the interests of hedgers, as the prices will be lower and the market more efficient.

The liquidity of the exchange-traded products is a very important benefit to investors. Unfortunately, it is difficult to develop a new market, and it remains to be seen whether significant liquidity will find its way to this market.

TRENDS AND EXPECTATIONS

Catastrophe derivatives and industry loss warranties occupy an important place in the universe of insurance-linked securities. They provide capital markets participants with a way to invest in insurance risk without having to worry about moral hazard, or potential inadequacy of the risk analysis due to a specific underwriting portfolio being significantly different from that of the insurance industry as a whole. The standardised nature of these instruments is a significant contributor to the potential overall growth of the ILS markets.

Projections of future developments are very difficult when it comes to

most individual types of insurance-linked securities. Longer-term qualitative forecasts happen to be easier than specific short-term predictions. When it comes to projections of long-term growth of the catastrophe derivatives and ILWs together, as one category of the insurance-linked securities market, they are very positive. While the exchange-traded products have the most growth potential, they have not yet reached the critical mass necessary to assure this growth. The ILW products, however, are past the point of any doubts related to their continuing existence, and they will continue to play an important role in the securitisation of catastrophe insurance risk.

These positive expectations, not fully conclusive, are based on the following observed conditions and trends.

- ❑ Any product standardisation makes it easier to attract investor capital. Index-linked products address a number of investor concerns and make it easier for new investors to enter the marketplace of securitised insurance risk.
- ❑ Development of exchange-traded products can bring liquidity to the market where buy-and-hold strategy is standard for investors. Hedgers are rarely concerned with future liquidity at the time they purchase protection. However, they too will benefit from it as liquidity will give them the option of dynamic hedging to provide the most efficient protection. The increase in liquidity would also tend to decrease price levels for these instruments.
- ❑ These products add to the toolbox available to an investor for effective assembling and optimisation of an ILS portfolio. They facilitate dynamic portfolio management and allow the investor to move further away from the less efficient buy-and-hold strategy.
- ❑ Growing transparency is beneficial not only to the exchange-traded segment of the catastrophe derivative and ILW market. Settlement prices for products such as IFEX event-linked futures are growing in importance as a reference point for the traditional catastrophe reinsurance market. As greater attention is paid to these products, they might be considered more often as a substitute for some layers of traditional catastrophe reinsurance.
- ❑ At least for IFEX contracts, there have been more than one market maker posting daily bids and offers for the most popular contracts. They provide some liquidity to the market and facilitate on-screen transactions. The existence of market makers for these contracts is one of the important ingredients for future growth of the market.

- ❑ Clearing block trades of catastrophe derivatives through the exchange essentially eliminates the credit-risk issue present in most types of insurance risk transfer.
- ❑ Exchange-traded products provide the most flexibility to quickly react to changing conditions. For example, they are perfect instruments, assuming sufficient liquidity, for live cat trading, where protection buyers can hedge their exposure in the face of an approaching hurricane; opportunistic investors can take advantage of the same situation.

Additional developments leading to market transformation and potential growth are likely to be based on the following factors.

- ❑ New indexes in addition to the ones mentioned above can enable transfer to the capital markets the risks that are currently residing almost entirely in the insurance and reinsurance industry. Political risk and aviation liability are two examples of such risks. Insurance and reinsurance capacity for these risks can be limited, leading to some risks remaining uninsured and forcing corporations to retain them even when it is not prudent.
- ❑ New parametric indexes (such as Paradex developed by RMS) simplify the transfer of insurance risk to the capital markets, and can facilitate the OTC insurance derivative transactions as well as other types of ILS.
- ❑ Catastrophe derivatives and ILWs are gaining broader recognition as sources of retrocessional capacity (even when not done in reinsurance form) at times when capacity levels are unstable and traditional capacity is clearly insufficient.
- ❑ Changes in accounting rules, though unlikely in the near future, may eliminate the accounting disadvantages for insurance and reinsurance companies of buying catastrophe protection in the derivative as opposed to traditional reinsurance form.
- ❑ The use of hurricane and other catastrophe derivatives as part of the comprehensive management of commodity investment portfolios can open up new markets for these products and contribute to greater market efficiency. Weather derivatives are already used for this purpose.

ILWs in the traditional reinsurance form, OTC catastrophe derivatives and exchange-traded derivatives all provide an efficient way for the transfer of catastrophe risk to capital markets. These instruments are expected to play a growing role in the insurance-linked securities markets, due to the unique advantages they provide to both buyers and sellers of catastrophe protection.

Reinsurance Sidecars and Securitised Reinsurance

SECURITISATION OF REINSURANCE

Investing in reinsurance companies, whether in the form of common stock, preferred shares or debt, has always attracted investors searching for companies that are undervalued and those best positioned for profitable growth. Of particular interest are the companies underwriting reinsurance lines of business, where capacity is tight and rates are consequentially “hard”. They are seen, usually justifiably, as the best profit generators in the short run, until the markets correct themselves and inflows of capital or other events solve the capacity problem.

In an ideal world, it would be possible to invest not in the securities of the reinsurance company as a whole, but in specific types of reinsurance business – those that are the most profitable at the moment – and exit the investments when these pockets of extra profitability disappear or move away. Many types of insurance-linked securities are intended to provide investors, at least to some degree, with this very opportunity. Catastrophe bonds are a good example of such securities. There are opportunities to invest in insurance risk on an even more granular basis. For example, collateralised reinsurance can be a way to invest in a specific reinsurance contract. For risks that require significant capacity not found on acceptable terms in the reinsurance market, collateralised reinsurance can provide investors with exposure to a desirable type of risk and allow them to compete directly against reinsurance companies. However, this type of investing requires significant reinsurance-underwriting expertise found only in a few dedicated (ILS) funds that effectively underwrite reinsurance. When we define securitised reinsurance in this very narrow away, excluding catastrophe bonds and similar instruments, the investor universe becomes very small.

A way to invest in reinsurance risk underwritten by a reinsurance company is through a reinsurance sidecar structure. In this case, the investor

would benefit from exposure to the currently profitable business without needing to fully understand reinsurance underwriting, since this function remains with the sidecar sponsor. For the investor, underwriting then becomes underwriting of the reinsurance underwriter rather than underwriting of reinsurance. Of course, other considerations also play an important role in the investment decision. This chapter examines reinsurance sidecars, looks at their structure and the advantages and disadvantages of their usage from both a sponsor and investor perspective.

REINSURANCE SIDECARS

A reinsurance sidecar is a limited-life special purpose vehicle that provides reinsurance companies with additional capacity while allowing investors to gain exposure to pure insurance risk. Several characteristics differentiate reinsurance sidecars from other types of insurance-linked securities (ILS) and from direct investing in reinsurance companies. Reinsurance sidecars allow us to share in the narrowly defined types of insurance risk and return of the sidecar sponsor (reinsurance company) without taking on the multitude of risks involved in operating a reinsurance company. They also provide investors with a clear and clean exit strategy. Since the structure is usually quota share reinsurance, the risk and reward of underwriting predefined (typically property catastrophe) reinsurance lines of business are shared by the sponsor and the investors.

Effectively, reinsurance sidecars provide what could be called “accordion capital”, which can increase or decrease depending on the needs of the sponsoring reinsurance company. The lack of reinsurance capacity, along with high property catastrophe rates, in the aftermath of the 2005 hurricane season created a situation whereby sidecars were the best vehicles for addressing the sponsor need for capital while providing the investor with attractive risk-adjusted returns. That scenario served as in impetus for the development of this market.

SIDECAR STRUCTURE

A simplified diagram of a reinsurance sidecar structure is presented in Figure 6.1. Reinsurance Company acts as the sponsor of Sidecar Re. Sidecar Re is the entity that enters into a reinsurance contract with Reinsurance Company. The reinsurance coverage is collateralised with proceeds from issuing securities (equity and debt or only equity) to investors and from reinsurance premiums. Reinsurance Company would usually participate in the equity tranche (primarily for psychological reasons, to provide extra

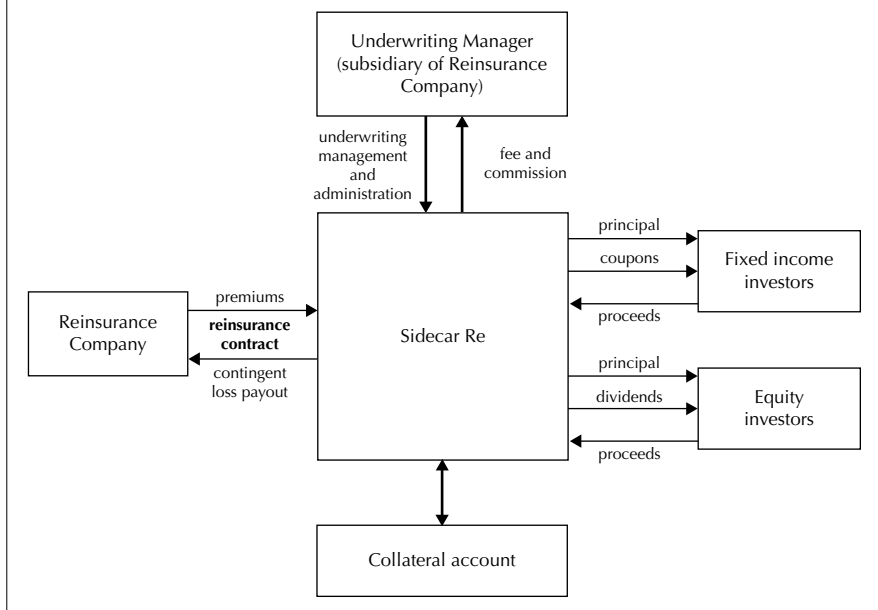
reassurance to other investors of the alignment of interest). Sidecar Re itself can be initially set up by Reinsurance Company or other parties such as an investment bank structuring the transaction.

Sidecars used to have a very heavy debt component, often with more than one debt tranche. The coupons can be fixed or, as in the typical case of catastrophe bonds, stated on the “Libor-plus” basis. The debt may be rated. As discussed later, the situation has changed and this type of leverage is rarely available in the current investment environment.

Sidecar Re can be a special-purpose reinsurance company or a combination of a holding company and an operating reinsurance company. A collateral account is set up as a trust with permitted investments and rules governing the release of collateral.

The specific structure illustrated in Figure 6.1 includes Underwriting Manager, a subsidiary of Reinsurance Company that is compensated by Sidecar Re for providing underwriting and management services; profit commission could be part of the compensation. The Underwriting Manager element is optional and could be excluded from the structure. Reinsurance Company can provide all these services and be compensated for them as part of the reinsurance agreement with Sidecar Re. This element is more

Figure 6.1 Sample structure of a reinsurance “debt and equity” sidecar



likely to be found when the sidecar assumes part of the exposure directly rather than through a reinsurance contract with the sponsor. Underwriting Manager does not necessarily have to be a subsidiary of Reinsurance Company; both could be subsidiaries of a common holding company. This element is not to be confused with that of a company operating sidecars. (Horseshoe Group is an example of such company.)

When a holding company structure is used, the holding company would issue equity but the debt would usually be issued by the operating reinsurance company. Fixed-income investors would typically have a security interest in the holding company, which in this case would be similar to parental guarantee.

A sidecar would normally not be rated; the rating is not required due to the collateralised nature of reinsurance protection involved. The collateral is usually held in a trust governed by New York Regulation 114.

If at the end of the exposure period there are loss reserves – both actual claim reserves and the reserves for incurred but not-yet-reported claims (IBNR) – the life of the sidecar is extended for a pre-agreed period of time, at the end of which any remaining liabilities are commuted according to the rules in the reinsurance agreement. In some cases, most of the funds could be returned to investors at the end of the exposure period, while the remainder – equal to the loss reserves and a safety margin – could be held longer until the liabilities are commuted. Since the great majority of sidecars reinsure short-tail property business, loss reserves are rarely an issue.

A sidecar can be fully capitalised from day one. Or, similar to the way it is done in the private-equity world, capital calls can be made periodically as the sidecar grows with new business being written.

INVESTOR PERSPECTIVE

Reinsurance sidecars offer investors the same advantage as most insurance-linked securities – exposure to an asset class with low correlation to the rest of the financial markets. In addition, there are advantages that are specific to investing in sidecars, the most important of which are the following.

- ❑ Investors in sidecars can be extremely opportunistic by investing where and when excess profits are expected to be generated. For example, if reinsurance capacity in property catastrophe lines is limited after a hurricane or another event, investing in a property catastrophe reinsurance sidecar might generate high returns. Some sidecars created after the 2005 hurricane season have already closed with returns in excess of

60% per annum for their equity investors. (At the same time, however, investors in some such sidecars lost money.)

- ❑ Reinsurance sidecars provide relatively pure investment in specific types of reinsurance risk. Unlike investing in securities issued by a reinsurance company, in this case investors have limited concern about profitability of other lines of business written by the company, loss reserving issues, or performance of the reinsurance company's investment portfolio. They invest strictly in the type of reinsurance risk they want, trying to take full advantage of its low correlation with other financial assets and of its potential for high profitability.
- ❑ Investors in reinsurance sidecars avoid the need to develop their own expertise in reinsurance underwriting by gaining access to the expertise of underwriting teams, which they are able to evaluate before making an investment. If properly structured, reinsurance sidecars provide an alignment of interest between the sponsors (reinsurance companies) and investors.
- ❑ The limited lifespan of a sidecar provides a clear and clean exit to investors. The money does not become tied up for an uncertain period of time, as the investment is made only for the time that the underwriting conditions are expected to remain favourable. If the reinsurance markets soften, investors simply will not reinvest and will choose to employ their capital elsewhere.
- ❑ In addition to having a clean and defined exit, sidecar investments are relatively easy to enter. An investor wishing to take advantage of favourable underwriting conditions (hard markets) in a specific line of business might otherwise need to go through the trouble of starting a new reinsurance company to focus on this market niche – with all of the time and trouble required to set up a new entity and assemble a strong management and underwriting team for the startup – and wait for the ramp-up of business. Alternatively, a private equity investor might decide to buy an existing reinsurance company, which carries the risk of inheriting legacy issues that are difficult to uncover in the due diligence process. Investing in a reinsurance sidecar avoids all of these complications.

SPONSOR PERSPECTIVE

For a reinsurance company, sponsoring a sidecar can be an efficient way to get access to capital to underwrite more business. The main advantages of sponsoring a reinsurance sidecar are the following:

- ❑ Sidecars offer a source of capital that is not dilutive to current shareholders while providing access to additional capacity to underwrite the line of business perceived to be the most profitable.
- ❑ Since sidecars have a limited lifespan, the capital is guaranteed to go away. This is an advantage compared with the situation where equity capital is raised by the sponsor or its holding company to take advantage of hard markets, but in two or three years, when the opportunity has exhausted itself, the sponsor has to face hard decisions on what to do with the extra capital. Artificial measures such as share buyback at unpredictable prices are not needed when the sidecar solution is utilised instead of raising traditional equity.
- ❑ In some cases, sponsoring a sidecar can be faster and can involve smaller expenses than issuing securities through the reinsurance company itself.
- ❑ Retrocessionary coverage becomes very expensive when capacity in a specific line of business is tight. Sponsoring a reinsurance sidecar can provide a cheaper alternative to retro coverage. When retro coverage is prohibitively expensive, the choice might be between sponsoring a sidecar and curtailing underwriting (instead of expanding it), as these are the times when underwriting is most profitable.
- ❑ Fee income to the sponsor in the form of ceding and profit commission can be significant, depending on the terms negotiated for the sidecar.
- ❑ The extra capacity resulting from sponsoring a sidecar can allow the reinsurance company not only to expand its underwriting but also to write larger lines (provide greater limits) and help its clients (cedents), who might otherwise need to split the limits across additional parties. In some cases, the ability to write larger lines opens new markets for the reinsurance company; these markets might not be open to smaller players.

These advantages by themselves do not imply that reinsurance sidecars are better than other instruments such as, for example, catastrophe bonds. Each has its advantages and disadvantages, and the choice is dictated by the specific situation of the sponsor and the market conditions at the time.

SIDECAR TYPES

There are a number of structures that fall under the general umbrella of sidecars. Sidecars can reinsure individual lines of business or a combination of several lines of business. Hannover Re was first to introduce sidecars that combine several types of risk in the same vehicle. In addition to property catastrophe reinsurance and other property risks, marine and aviation lines

lend themselves particularly well to sidecar-type investments. These are both short-tail lines with claims settling fast and allowing for short tenor of sidecar securities. They also tend to suffer periodically from capacity crunches that sidecars are designed to alleviate. Life insurance risks have also been transferred to the capital markets using the sidecar approach.

Almost every sidecar has a quota share reinsurance agreement, but an alternative structure can instead have an excess-of-loss reinsurance contract between the sponsor and the sidecar. The difficulty with excess-of-loss structures is that they lack the direct alignment of interests between the sidecar sponsor and the investors that comes more naturally in a quota share type of reinsurance agreement.

Most sidecars function in a straightforward manner in that they reinsure a certain percentage of the sponsor's business that meets specific guidelines. There is only one reinsurance agreement – that with the sponsoring reinsurance company. Sometimes another approach is utilised, where the sidecar assumes business directly from cedents in proportion to the business being assumed by the sponsor. The underwriting is still performed entirely by the sponsor, and the sidecar fully relies on the sponsor in this regard, as well as when it comes to administration issues. Each cedent then has agreements with two parties, the sponsoring reinsurance company and the sidecar. The first structure, which is the more straightforward and involves only one reinsurance contract, is the better one to implement unless there are special circumstances that make the other structure more attractive.

Leveraged versus equity-only sidecars

The sidecars that came to existence after the devastating hurricane season of 2005 have practically all been leveraged, often quite heavily. For example, the Panther Re sidecar sponsored by Hiscox (see Table 6.1) consisted of US\$144 million in equity and two debt tranches with different ratings, of US\$72 and US\$144 million, for a total of US\$360 million in capital.

Debt tranches of a sidecar, especially the higher ones, are somewhat comparable to catastrophe bonds. The probability of default is very low, unlike the probability that the equity tranche will suffer losses. An argument can be made that the yield should also be comparable, with some extra yield to account for the fact that the probability of default can be estimated only with a rather lower degree of accuracy compared with the analysis performed for catastrophe bonds. However, this was not always the case when debt funding of sidecars was common after the 2005 hurricane season. There is a perception that some debt tranches may have been underpriced,

and that they were purchased by investors who did not have the expertise to value them properly. (This statement does not refer to any specific transaction mentioned.)

After the Lehman bankruptcy and the general credit crisis, leverage became unavailable or very expensive. Demand from investors in the debt tranches of reinsurance sidecars all but disappeared, at least on the terms offered. When the next generation of sidecars tried to raise money starting with the autumn of 2008, at the time that the property catastrophe reinsurance rates increased and the capacity was limited, the investor community showed interest only in the equity-only sidecars. Even these sidecars for the most part were unable to raise funds, since investors demanded returns that appeared unreasonable to potential sponsors. It is unclear whether the leveraged sidecar will ever return; the only structure used in the future may be equity-only sidecars.

Representative sidecar transactions

Table 6.1 shows a partial list of sidecars issued in recent years. While the sidecar era officially began after the 2005 Katrina–Rita–Wilma hurricane season (the term “sidecar” became widely used only in 2006), the first sidecar transactions were performed before that and the structure by itself was already known. The first sidecar listed, Top Layer Re, was put together in 1999 to provide property catastrophe risk transfer to the capital markets starting in 2000. However, there were sidecar-type transactions even before that.

Two names, RenaissanceRe and Hannover Re, stand apart as pioneers of this type of risk transfer. While the focus of RenaissanceRe has always been on property catastrophe risk, Hannover Re has transferred to the capital markets risks ranging from property catastrophe to life insurance.

Table 6.1 is intended to serve only illustrative purposes. The list of sidecar transactions is much longer. There have also been a number of private deals that have never been publicly disclosed but utilised the same types of structures.

INVESTOR UNIVERSE

For reinsurance sidecars, the investor universe differs from that for catastrophe bonds and most other types of insurance-linked securities. Dedicated ILS funds are an important part of this universe, but their share has been well below what might be expected.

The equity tranches of sidecars attract private equity investors that generally do not participate in the ILS markets. The returns for the 2005–2007 class

Table 6.1 Representative list of reinsurance sidecar transactions

First year of coverage	Principal sponsor(s)	Reinsurance sidecar	Line of reinsurance business	Initial size as reported (US\$ millions)
2000	Renaissance Re and State Farm	Top Layer Re	High excess of loss US property catastrophe	100
2002	White Mountains Re	Olympus Re	Combination of property catastrophe, marine and retro	500
2006	Montpelier Re	Blue Ocean Re	Property catastrophe retro	300
2006	XL Capital	Cyrus Re	Property catastrophe regular and retro	525
2006	Arch Re	Flatiron Re	Property reinsurance (mostly catastrophe) and marine	900
2006	Renaissance Re	Starbound Re	Property and marine reinsurance	310
2007	Hiscox	Panther Re	Property catastrophe reinsurance	360
2007	ACE	MaRI	Property catastrophe reinsurance	400
2008	Hannover Re	Globe Re	Property catastrophe retro	133
2009	Hiscox	Syndicate 6104	Property catastrophe reinsurance	62
2009	Renaissance Re	Timicuan Re II	Reinstatement premium protection for US property catastrophe	60

Notes: Year of inception may be earlier than the first year of coverage shown above. Only the initial Cyrus Re transaction (inception in 2006 providing coverage for 2007) is shown. The additional raise (US\$100 million) for Cyrus Re is not included. Transaction size includes equity investments, if any, by the sponsor.

of sidecars were sufficient for this class of investors. The relatively short lifespan of sidecars was another advantage, as it provided the exit often elusive in private equity investing. Some of the investors otherwise might have considered setting up reinsurance startups to take advantage of the hardening property catastrophe rates. Sidecar equity investments gave them a more efficient way to achieve the same goal.

Some of the investors in debt tranches were new to the insurance and reinsurance space and made their decisions based primarily on ratings or, for unrated tranches, on rather limited analysis. That is the way investors often enter a new market; later, they gain greater expertise and make more informed decisions. In this case, however, there is a chance that sidecar debt tranches might never again become viable investments.

CONSIDERATIONS IN INVESTMENT ANALYSIS

Analysis involved in making investment decisions regarding sidecars includes all the same considerations as are present in analysing catastrophe bonds and similar types of insurance-linked securities. The catastrophe modelling analysis should be carefully examined and modified if necessary. Model choice and assumptions are part of this examination. However, there are also some important differences.

Expected rate of return is calculated both on the deterministic basis and, to the degree possible, on the stochastic basis. The deal cash model needed for this calculation is built based on the parameters of the transaction and a large number of assumptions. The uncertainty involved is significantly greater than in the modelling of catastrophe bonds. It is possible to utilise modelling software for property catastrophe risk, but the assumptions needed to be made about the risk exposure are usually so broad that in some cases we could wonder exactly how much value is added by this analysis. Coming up with a number of scenarios that might or might not be based on the output of the catastrophe modelling software, and assigning probabilities to each of the scenarios, is an approach that can produce reasonable results if done properly. Assumptions to be made in any analysis are numerous, and many of them have to do with the quality of the underwriting team of the sponsor. In addition, the analysis requires making assumptions about future market conditions over the lifespan of the sidecar; more than one scenario might be required. Prior performance of the underwriting team is very important, but it is also important to recognise that the new conditions can affect the behaviour of the underwriting team, including its underwriting standards and the level of risk aversion.

It is essential to examine the sidecar structure and understand to what degree the interests of the sponsor and investors in the sidecar are aligned. The type of business being reinsured should be strictly defined. The way profit commission is determined is of particular significance, as it might lead to misalignment of interests between the two parties. A similar issue concerns who bears the expenses of the deal, and how expenses are calculated, including also the ongoing expenses. Other ongoing concerns have to be addressed, including the need for and cost of monitoring to ensure that the risks transferred to the sidecar are within the parameters of the reinsurance agreement, and that the agreed-upon underwriting guidelines are being followed. Other reinsurance inuring to the benefit of the parties should be taken into account. Review of the documentation should cover areas such as handling of collateral and rules for releasing funds from the trust account; commutation, which is a significant point in the deal timeline; procedures for reserves valuation; and many others. Legal and accounting issues have to be analysed as well. Compliance with regulatory requirements, including reporting requirements, and assuring that the sidecar maintains tax-exempt status in its jurisdiction, are at the top of the list.

When an investor cannot become entirely comfortable with some of these important elements, the choice is either not to enter into the transaction or to require higher returns in recognition of the additional risk.

There is a difference between sophisticated models and those that are simply complicated. In the sidecar analysis, where so many assumptions have to be made, simpler approaches often work best. It can be easy to create a very complicated model based on numerous assumptions; it is also possible that this model might have little to do with reality and might be inferior to much simpler analysis. Sensitivity testing, always important in investment analysis and pricing, is especially important in this case.

The analysis of sidecars can be particularly difficult in the context of portfolio management of insurance-linked securities. The need to make numerous assumptions for these specific securities, along with their likely high degree of correlation with other property catastrophe insurance-linked securities, makes optimisation of an ILS portfolio that includes sidecar investments particularly challenging.

TRENDS AND EXPECTATIONS

Securitized reinsurance, in one form or another, is growing in importance in the world of reinsurance and insurance risk transfer. While most insurance-linked securities can be considered, at least to some degree, to be a form of

securitised reinsurance, the more granular approach, such as direct investing in a specific reinsurance contract, has been growing in popularity. This is relevant to the investment funds that have built in-house reinsurance underwriting expertise. This more narrowly defined securitised reinsurance is usually provided in the fully collateralised form.

Reinsurance sidecars, on the other hand, present a way to invest in a profitable underwriting risk without the need to have extensive reinsurance underwriting expertise. In fact, it is an inexpensive way to get access to top-level underwriting expertise. The key advantages of sidecars are the following.

- ❑ Investors gain exposure to reinsurance risk, with its low correlation with traditional financial markets, at the time and for the types of risk with the highest expected profitability. Reinsurance sidecars provide both an easy entry and a clear time-defined exit. The exit strategy does not need to be worked on: the exit is automatic.
- ❑ Reinsurance companies sponsoring sidecars gain immediate access to extra capital, allowing them to expand the underwriting activities in the lines of business that are considered to be the most profitable. Sidecars avoid the need for the reinsurance company to raise equity that would be dilutive to existing shareholders and might create complications later when the capital is no longer needed.

The future of reinsurance sidecars is uncertain. On the one hand, this type of insurance-linked security has proved to be very useful, as in the aftermath of the Katrina–Rita–Wilma 2005 hurricane season, when capacity in reinsurance markets was limited and additional capital was required. On the other hand, issuance of new sidecar investments has all but stopped, and other types of ILS and reinsurance have proved to be, in most cases, better alternatives to sidecars in the current environment. In 2008 and 2009, several sidecars were offered to or discussed with investors, but few were actually issued.

Since the reinsurance sidecar market has experienced real turmoil, and uncertainty persists, the following observations and trends are relevant.

- ❑ The current trend has turned against sidecars. In part, this has happened because sidecars can no longer issue cheap debt to investors and use it to provide leverage to sidecar equity investors. Potential investors in sidecar debt have less interest in the relatively low yields provided because they

cannot use leverage to augment their return. Leverage has become either unobtainable or very expensive, in stark contrast with the time after Hurricane Katrina, when most sidecars were established.

- ❑ The only viable sidecar structure in such conditions is an equity-only sidecar; issuing sidecars with debt tranches does not attract investor interest.
- ❑ Even equity-only sidecars have recently found it difficult to attract investors, who seek returns in excess of 20% or even 30% with relatively low volatility and are often not convinced of the quality of the underwriting team. In the absence of leverage provided by debt tranches, it is more difficult to achieve high returns on equity.
- ❑ Innovation in the way sidecars are structured and what type of risk is transferred to investors can increase the appeal of this instrument. Sidecars could be created for lines of business that have traditionally been viewed as not well suited for such instruments, such as longer-duration insurance; it could be transferred to investors through a sidecar if effective commutation mechanisms are developed. Difficult-to-model risks, including that of manmade catastrophes, present another example.
- ❑ A combination of several lines in the same sidecar – similar to the type that has been utilised by Hannover Re for years in some of its “K” transactions – might grow in volume, though such a solution would always have appeal only to a very small group of investors.
- ❑ There is some limited liquidity in the sense that, unlike in the case of reinsurance contracts, an investor can usually exit a sidecar investment early, even though it would probably not be on the most attractive terms. Brokers who provide indicative pricing for and facilitate trading of catastrophe bonds have sometimes also brokered sidecar transactions in the secondary market and even provided indicative prices for some tranches. The ability to exit sidecar investments early (even though they have short tenor from the very beginning) is important for some investors.

Sidecars have proved themselves to be an efficient way for investors to gain exposure to some of the most promising types of reinsurance risk, while at the same time providing reinsurance companies with additional capital when it is most needed.

Sidecars turn reinsurance companies into “accordion” reinsurers by providing them with capital when it is needed. One should not judge the success of sidecars by the year-on-year change in issuance. They are designed to be either more or less attractive depending on market conditions. The

marked slowdown in the appearance of new, or renewal of existing, sidecars that followed the explosion in sidecar issuance that happened after the 2005 hurricane season is actually indicative of this advantage of sidecars, as opposed to being a negative reflection on sidecars as an asset class. Sponsors have a choice of instruments in their toolbox, and they can use the ones most appropriate for market conditions at any given time. Reinsurance sidecars are one such instrument at their disposal, which they can use to advantage when the opportunity presents itself.

Credit Risk in Catastrophe Bonds and Other ILS

CREDIT RISK

Until 2008, and in particular the Lehman bankruptcy, credit risk was generally below the radar screen for sponsors, structurers and investors in insurance-linked securities. Since then credit issues have put a complete stop to some ILS transactions and forced structural changes to others. Credit risk emerged as an important issue in structuring insurance-linked securities. This chapter analyses the credit risk embedded in ILS with a particular focus on catastrophe bonds. It also provides an overview of the emerging solutions to mitigating credit risk in structuring these securities.

In a financial transaction between two parties, credit risk is the risk of a counterparty's default on its obligations, whether in whole or in part. The default can be in the form of nonpayment or payment reduced relative to the agreement; untimely payment; reduction in the obligor's credit ratings; or failure to maintain assets in an account at an agreed level or of an agreed quality. The definition depends on the legal agreement and typically does not include all of these elements.

Present in virtually all financial transactions, credit risk is a fact of life in the world of finance. Credit risk of a transaction has to be analysed and quantified; the results are then incorporated in pricing the transaction. An investor wants to be compensated for any risk, including credit risk. In the case of fixed income securities, where credit risk is the risk that drives their performance, it is the primary determinant of the yield these securities can command in the market. Catastrophe bonds, though structured as fixed-income securities, are not supposed to have significant credit risk and instead are intended to be a vehicle for transferring catastrophic insurance risk. Credit risk in such transactions affects their sponsors as well; the risk is not limited to investors.

Ways of mitigating credit risk

A number of ways to manage and reduce credit risk have been developed. Some of them could not only reduce the probability of a credit event but also attempt to reduce its negative effect if it happens. These include the following.

- ❑ Collateral is the most common way to reduce credit risk in a transaction. Assets or rights to the assets (as defined and with restrictions stipulated in accompanying legal documents) are pledged to a party (in the most general case, a lender) or deposited in a separate account (usually a trust) for the party's benefit. A collateral account is an efficient way to provide protection against credit risk, in particular when its assets are liquid and properly valued. In ILS, the collateral account can also play a broader role than protection against credit risk.
- ❑ Overcollateralisation is a way to avoid the risk that the value of the assets will decrease or that liquidity concerns will make their quick sale impossible unless done at lower prices.
- ❑ A guarantee (loan guarantee in the traditional credit world), such as a guarantee provided by a parent company, reduces credit risk as well, especially if the guarantee is unconditional. Parental guarantees have been used in a number of ILS transactions, in particular on the life insurance side. They might have unintended consequences such as change in ratings of the parent; they are less valuable than guarantees provided by an unrelated party.
- ❑ Letter of credit (LoC) is another way to reduce credit risk. It can be in the form of a guarantee issued by a bank or another financial institution. In some cases, it can serve as a substitute for a collateral account. To reduce credit risk, an LoC has to be irrevocable.
- ❑ Credit derivatives have been used to mitigate credit risk. This is a less popular method of credit-risk mitigation, especially nowadays, and has less relevance to insurance-linked securities.
- ❑ Credit insurance as a form of credit-risk mitigation is of relevance to ILS when it is in the form of a financial guarantee provided by a monoline financial guarantor. This type of credit enhancement has enabled many ILS transactions but now it is generally unavailable due to the financial difficulties of every single financial-guarantee company. It is likely that, in the future, credit wrap provided by a monoline financial guarantor will again be used in some of the ILS structures, but on a much smaller scale than in the past. Financial guarantee to enable a Regulation XXX-type transaction is unlikely to be available.

There are also other ways to mitigate credit risk, but not all of them are applicable to the credit risk of insurance risk securities.

CREDIT RISK AND ILS

Catastrophe bonds and many other insurance-linked securities have been conceived and structured as a way to help investors gain exposure to pure insurance risk, with all other types of risk stripped away to the degree that they would be negligible. At the very least, that was the goal. Pure insurance risk provides diversification to investors due to its low correlation with the traditional financial assets.

Credit risk was supposed to be almost absent in investing in securities such as catastrophe bonds. It was supposed to be relatively low even in unwrapped XXX securities that transferred redundant life insurance reserves to the capital markets. There was no perceived need to change anything in the ILS structures to reduce credit risk more than was already the case.

In property catastrophe bonds, the investment analysis included examination of credit risk only to the degree that the legal documents were conforming to the established standard and one of the traditional counterparties was providing the total return swap. Once the legal documents were judged to conform to the standard, the “real” analysis started, with its exclusive focus on modelling the risk of natural catastrophes and the insured losses resulting from the catastrophic events. Credit risk was considered to be negligible relative to the “real” risk of catastrophe bonds, the risk of insured losses due to a hurricane or an earthquake.

TRADITIONAL SOLUTIONS

The credit risk issue was not neglected in the past: rather, it was analysed and then considered to have been fully addressed in the standard structures used for catastrophe bonds and other types of insurance-linked securities.

Cat bonds

The standard structure of a catastrophe bond is described in Chapter 3. The elements of the structure intended to manage credit risk are the collateral account (trust) and the swap counterparty. The total return swap was initially introduced primarily for the purpose of eliminating interest rate risk from the transaction; the significance of credit risk was not fully appreciated at the time. In the traditional cat bond structures used until 2009, returns from the collateral account were swapped for a Libor-based rate with a highly rated counterparty. The counterparty rating, which sometimes

later deteriorated, was considered to be sufficient protection against credit risk, and the rules on permitted investments in the collateral account were not particularly strict. In addition, there was no mechanism for independent valuation of the collateral account assets on a frequent basis that would also require immediately adding assets should the collateral value fall below a certain level. The existence of the total-return swap arrangement was viewed by most to be a sufficient guarantee that credit risk was not an important issue in catastrophe bonds. That view was proved to be wrong.

Securitisation of Regulation XXX reserves

Typical securitisation structures for funding Regulation XXX life insurance reserves have included two components that need to be considered from the point of view of credit risk. First, they have a collateral account, typically in the form of a Regulation 114 reinsurance trust. Such trusts used to be considered to be extremely secure; this is no longer the majority view. Second, a typical element was the financial guarantee of the type provided by a mono-line financial guarantee company such as AMBAC or MBIA. The financial guarantee was used to enhance the ratings of securities offered to investors. The AAA ratings significantly expanded the universe of potential investors and made the securities liquid despite their very long tenor. As the financial guarantors lost their high ratings, so did these securities.

Extreme mortality bonds

Extreme mortality bonds have a potential weakness similar to that of property catastrophe bonds in that there is a reliance on a swap counterparty and insufficient guidelines and controls for the management of a collateral account. In addition, the credit wrap provided for some tranches by mono-line financial guarantee companies has suffered from the weaknesses described above.

THE NEED FOR NEW SOLUTIONS

The issue of credit risk in catastrophe bonds was brought to light very quickly when Lehman declared bankruptcy. The credit risk in cat bonds – something completely disregarded by the vast majority of investors and other parties involved – suddenly became such a significant issue that the very structure of cat bonds was questioned and new issuance completely stopped.

Cat bonds with Lehman as TRS provider

Four catastrophe bonds had Lehman as their TRS provider at the time when Lehman went bankrupt in 2008. They were quickly downgraded by rating agencies. The four were:

- ❑ Ajax Re Ltd's Class A Series 1 (sponsored by Aspen Insurance Ltd);
- ❑ Carillon Ltd's Class A-1 (sponsored by Munich Re);
- ❑ Newton Re Ltd's Class A 2008-1 (sponsored by Catlin); and
- ❑ Willow Re Ltd's Class B 2007-1 (sponsored by Allstate).

Initially, there was some hope that a replacement TRS counterparty would be found. As the problems with assets in the collateral accounts became apparent, these hopes were dashed. Defaults followed.

The consequences of Lehman's bankruptcy put a stop to new issuance of catastrophe bonds and led to the re-evaluation of the risk embedded in bonds that had been issued but not yet retired. Investors demanded new solutions in order to become comfortable with the credit risk of cat bonds. New problems came to the surface, such as the difficulty of getting information on what assets were in the collateral accounts. The element of transparency was clearly missing in most of the cat bond transactions.

Financial guarantee

When financial guarantee suddenly became unavailable as the monoline financial guarantors were downgraded, this type of credit enhancement ceased to be an option for new issues, dramatically changing the securitisation landscape for the types of ILS that needed this kind of credit enhancement to attract investors. This change also wreaked havoc for the owners of wrapped securities that were suddenly downgraded and, in some cases, became illiquid. It is interesting to note that some such securities happened to be in the collateral accounts of catastrophe bonds; their sudden downgrade and lack of willing buyers contributed to the predicament. Here too new solutions were needed.

SOLUTIONS TO CREDIT RISK ISSUES IN INSURANCE-LINKED SECURITIES

There is no simple solution to the lack of financial guarantee provided by monoline financial guarantee companies. Even if it becomes available again, the cost is likely to be prohibitive for these transactions.

For credit issues involving catastrophe bonds, however, several solutions have emerged. The issuance of catastrophe bonds resumed at the beginning of 2009, with more than one solution being utilised.

The main collateral solutions that have been either proposed or utilised in structuring these "new and improved" catastrophe bonds include the following:

- ❑ Bank deposits/CDs with highly rated banks are an easy solution to implement, addressing credit risk issues with the exception of the counterparty (bank) default. Such deposits are unsecured, and the bank has to be rated AAA or AA to make this solution acceptable.
- ❑ Government-guaranteed bank debt – in the form of the US Temporary Liquidity Guarantee Program (TLGP) – with TRS presents another solution. The weakness of this solution is that the FDIC guarantee expires on December 31, 2012 (unless extended again). The cat bonds that have utilised this collateral solutions include:
 - Atlas V Series 1, 2 and 3 (sponsored by SCOR);
 - East Lane Re III Series 2009–1 Class A (sponsored by Chubb);
 - Mystic Re II Series 2009–1 (sponsored by Liberty Mutual); and
 - Ibis Re Classes A and B (sponsored by Assurant).
- ❑ US Treasury money market funds represent another solution that is simple and effectively eliminates credit risk. The only problem it presents is the very low rate of interest on these securities, which makes it more expensive for the sponsor to provide the yield required by investors. The cat bonds that have utilised this solution include:
 - Successor II Series 4 Class F (sponsored by Swiss Re);
 - Residential Re 2009 Classes 1, 2 and 4 (sponsored by USAA);
 - Parkton Re (sponsored by Swiss Re on behalf of NCJU/NCIUA);
 - Multicat Mexico 2009 Classes A, B, C and D;
 - Redwood Capital XI (sponsored by Swiss Re);
 - Successor X Series 2009–1 Classes I-U1, I-S1 and I-X1 (sponsored by Swiss Re);
 - Longpoint Re II Classes A and B (sponsored by Travelers);
 - Lakeside Re II Class A (sponsored by Zurich); and
 - Foundation Re III Series 2010–1 Class A (sponsored by Hartford).
- ❑ Triparty daily repurchase structure goes a long way towards minimising credit risk, but it too has some disadvantages. Besides the seeming complexity of this approach, the credit risk of the repurchase counterparty might be correlated with the credit risk of the assets in the collateral account. The approach is not as complex, however, as it might appear to those unfamiliar with repos. The cat bonds that have utilised this solution include:
 - Eurur II Series 2009–1 Class A (sponsored by Hannover Re);
 - Montana Re Series 2009–1 Classes A and B (sponsored by Flagstone Re); and
 - Atlas VI Series 2009–1 Class A (sponsored by SCOR).

- ❑ Customised puttable notes issued by sovereign or quasi-governmental entities are another innovative solution that has been successfully used. AAA-rated notes issued by Kreditanstalt für Wiederaufbau (KfW) and the International Bank for Reconstruction and Development (EBRD) have been used as collateral in catastrophe bond structures. The cat bonds that have utilised this solution include:
- Blue Fin Series 2 Class A (sponsored by Allianz);
 - Ianus Capital (sponsored by Munich Re); and
 - Calabash Re III Classes A and B (sponsored by Swiss Re transferring a portion reinsured risk of ACE).

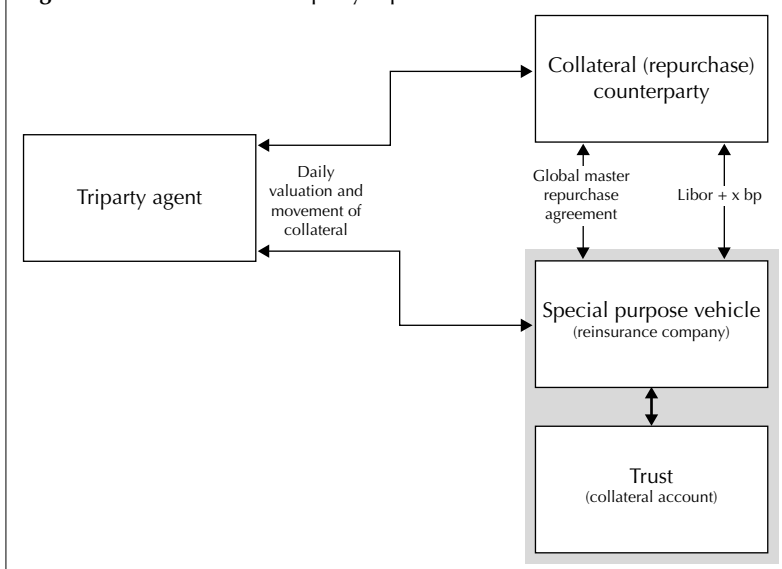
Some of these solutions are explained in more detail below.

TRIPARTY REPO ARRANGEMENT

One of the proposed solutions to the TRS issue in catastrophe bonds is the triparty repurchase agreement. In this arrangement, first used in structuring a catastrophe bond by BNP Paribas, the issuer of the bond, a bank (or another financial institution) and a third party taking on the role of a triparty agent enter into a repo agreement. Standard legal documentation is used for the agreement. Triparty repo agreements are not unique to the catastrophe bond market and have been used in financial transactions for many years.

First, eligible collateral is defined; it would typically include highly rated securities with high liquidity and easily observable prices. These would usually be investment-grade bonds that are sufficiently liquid. The term “eligible collateral profile” is sometimes used to describe permitted collateral composition. In a true-sale agreement, the collateral purchase counterparty (typically a bank with a sufficiently high rating) enters into an agreement with the issuer to sell it a pool of such eligible collateral investments in return for cash; the agreement calls for the purchase counterparty to repurchase the collateral at the end of a specified time period. The repurchase counterparty pays the issuer quarterly interest that is typically equal to the three-month Libor plus a spread (repurchase spread). The agreement also stipulates that a predefined level of overcollateralisation be maintained. The general outline of the overall structure is shown in Figure 7.1.

A key role is played by the triparty agent who manages the transaction. The triparty agent has to be completely independent. In the type of repo agreement discussed, the triparty agent’s primary functions are providing daily valuation of the collateral, assuring that the eligible securities rules established for the transaction are maintained, and managing daily move-

Figure 7.1 Outline of the triparty repo structure

ment of the investments between the parties based on the above-mentioned valuation. The eligible securities are chosen to be liquid, so the valuation is always on the mark-to-market basis. The whole procedure for the daily valuation and movement of the collateral is automated with little human involvement. The limited need for human involvement minimises operational risk that would otherwise be present when daily transactions are executed. The triparty agent also provides daily reporting services on the collateral composition, movements and substitution. These reports are received by the repurchase counterparty and the cat bond issuer; in a structure that provides greater transparency, access to daily reports can also be granted to the investors in the cat bond. The structure itself does not preclude the option of providing this information to all potential investors as opposed to only those who already own the cat bond securities.

Eligible collateral typically carries some restrictions in addition to the ones mentioned above. These could include: concentration limits; specific exclusion of some types of securities (such as mortgage-backed securities, CDOs or asset-backed securities) regardless of their rating; limits on the stated maturity of the securities; exclusion of all or some convertible instruments; exclusion of any securities issued or guaranteed by the sponsor, the repurchase counterparty, their subsidiaries or holding companies (excluded issuers); and requirements that the securities be denominated in a specific

currency (US dollar or euro). Excluded issuers may also include those whose credit ratings are perceived to have very high correlation with the rating of the repurchase counterparty.

The “adjustment percentages” are applied to the collateral investments if the repurchase counterparty credit rating deteriorates. The percentages differ by asset class and credit ratings. Adjustment percentages might also differ by the stated maturity of securities in the collateral.

It is important to emphasise that, in the repo structure, the cat bond issuer fully owns the collateral, which provides comfort if the repurchase counterparty defaults.

Legal considerations

All of the above is governed by the Global Master Repurchase Agreement (GMRA) entered into by the cat bond issuer and the repurchase counterparty. Under the GMRA, the cat bond issuer appoints a custodian to act as its agent in respect to entry into repurchase transactions under the GMRA. The repurchase counterparty, the custodian (on behalf of the issuer) and the triparty agent enter into an agreement that the triparty agent will act as an agent for the issuer and the repurchase counterparty. The repurchase agreement serves the purposes outlined above; the proceeds from the sale of cat bond securities are used to generate Libor- or Euribor-linked return collateralised by assets meeting the eligibility requirements described above and at the level of collateralisation specified. The triparty agent, through daily margining, provides for the evaluation of the collateral and the necessary movement of assets to meet all the collateral requirements. Conditions for the termination of the GMRA are clearly defined. (This clarity is important because some of the standard collateral documentation turned out to be inadequate when Lehman defaulted, leading to extra scrutiny of the termination conditions in all collateral documents.) In general, if the agreement is terminated due to the default of the repurchase counterparty, assets in the collateral accounts are liquidated and invested into predetermined types of assets (typically money market funds). Any amounts owed to the repurchase counterparty are paid first. The agreement may call for, or allow for trying to find, a replacement counterparty; there would be a time limit on the attempts to find such a replacement. A number of additional legal considerations lie outside the scope of this chapter.

The transaction is structured in such a way that the investors in the cat bond notes would not be in violation of the US ERISA rules by making the repurchase counterparty or the issuer a pension plan fiduciary. This could

happen if one of them is involved in “active management” of collateral investments.

CUSTOMISED PUTTABLE NOTES

Customised puttable notes provide another solution to the credit risk problem in catastrophe bonds. They avoid the need to use a total return swap (TRS) and the credit risk associated with the TRS counterparty. Highly rated (AAA) bonds issued by an entity that is backed by a government can be used as collateral. These customised puttable notes are designed to match the tenor of the cat bond securities. In the first transaction of this kind, Blue Fin sponsored by Allianz, the collateral was composed of puttable notes issued by Kreditanstalt für Wiederaufbau (KfW), the German Development Bank. The notes pay a Libor-linked return. They are puttable to Kreditanstalt für Wiederaufbau at the option of the holder after a certain period of time. The Ianus Capital bond sponsored by Munich Re also used customised puttable notes issued by Kreditanstalt für Wiederaufbau to provide extremely solid collateral without the need for a total return swap counterparty.

Another transaction of this type was Calabash Re III, a cat bond sponsored by Swiss Re. The collateral in this case is composed of medium-term custom notes issued by the International Bank for Reconstruction and Development (IBRD), which, like KfW, is AAA-rated.

Such custom notes are issued by AAA-rated entities that also have governmental backing. They are unsecured and unsubordinated. These notes generally rank equally among themselves as well as with all other securities and obligations of the issuer that are unsecured and unsubordinated.

The only potential disadvantage of customised puttable notes is that they have to be specially designed by an entity such as IBRD or KfW. It is easier to use, as collateral, investments that are readily available in the market. The need for customisation implies the necessity of a partnership with the government entity; and it takes longer to arrange such a solution.

USE OF US TREASURY MONEY MARKET FUNDS AS COLLATERAL

The use of Treasury money market funds is an effective solution to minimise credit risk in the collateral. The permitted investments in this case would be onshore US Treasury money market funds, with ratings of AAAm-G by Standard & Poor's or an equivalent by another rating agency. If such assets are for some reason unavailable, the structure might allow the use of federal money market funds that invest only in the obligations guaranteed by the

government or government agency and have AAAm rating by Standard & Poor's or an equivalent rating from another rating agency. Cash can always be the last resort. There is no TRS in this structure, even though this element can be included if necessary.

The disadvantage of this approach is the lower rate of return on these investments.

Relative importance of Libor-linked returns

Performance of some investors is judged on their excess return over Libor. If the Libor-linked returns are no longer offered in catastrophe bonds, these investors are at a disadvantage. Most investors want to achieve higher returns. Ultimately, it becomes a question of whether investors are willing to accept slightly lower returns in exchange for reduction in credit risk; or whether higher returns are more important to them, and they can live with some credit risk in their cat bond investments. Those who feel that cat bond investments should be true diversifiers with almost no correlation to other financial assets would choose to accept lower returns, since the credit risk is then almost completely eliminated.

COLLATERAL OPTIONS IN COLLATERALISED REINSURANCE

Collateralised reinsurance – in the narrow sense of providing collateralised reinsurance coverage for an individual reinsurance contract or an ILW in reinsurance form – has never featured uniformity in the choice of collateral solutions. Regulation 114 trust has been considered to be the norm, but many other solutions have been used as well. Obviously, within the Regulation 114 trust arrangement there are numerous options too.

In some cases, the collateral has been required to be posted only during the hurricane season. In others, partial collateralisation (not full limits) has been considered to be sufficient.

Even though uniformity in collateral arrangements is still missing from the narrowly defined collateralised reinsurance, the credit crisis and the Lehman Brothers debacle have brought attention to the credit risk of these arrangements. Most of the marketplace – but not all – has started paying very close attention to the collateral arrangements, leading to much stricter guidelines regarding eligible investments and concentration limits in the collateral account. In a few cases, the permitted investments were defined as federal money market funds only – representing a very high degree of conservatism never before applied to such arrangements. The overall movement to reducing credit risk in collateralised reinsurance continues.

TRENDS AND EXPECTATIONS

Credit risk in insurance-linked securities remains an area of uncertainty. Some solutions have disappeared with no replacement possible; for others a number of competing alternatives have been proposed and utilised, with no solution emerging that is clearly superior.

In the case of insurance securitisations that relied on credit enhancement provided by a financial guarantor, solutions are limited. They include the following.

- ❑ Tranches with higher attachment points can be introduced, where the rating will be sufficiently high (albeit not close to AAA) to attract a broader universe of investors. This is only a partial solution, however, since the issue of the lower-rated tranches is not addressed.
- ❑ Providing more information to investors would enable them to perform their own analysis and avoid overreliance on ratings. This solution focuses more on the lower-rated or unrated tranches that have greater risk. It also requires more investor education to enable them to better understand the risks involved and to develop greater confidence in the deal cashflow models.
- ❑ Another solution is to use alternatives to the full securitisation approach. These might include doing private deals or deciding to retain the risk and manage it in another fashion.

For insurance-linked securities such as catastrophe bonds, the fact that Lehman was the credit default swap counterparty for four such securities brought the credit risk issue to the surface in a manner that shocked most of the cat bond investors. The old structure became unacceptable and it became obvious that new solutions were required.

- ❑ Four collateral solutions for catastrophe bonds have been proposed, of which four have been used in the new generation of bonds that have been issued since the beginning of 2009: (1) using total return swap structures but with collateral invested in the FDIC-guaranteed securities; (2) utilising customised puttable notes issued by government-sponsored entities as collateral; (3) using the triparty repurchase agreement structure to minimise credit risk; and (4) investing the collateral in US Treasury money market funds. The last two solutions have been seen as the most promising, but the first two have not been rejected.
- ❑ At this point, it is unclear whether the dominant credit risk solution for

catastrophe bond securities will emerge in the near future. It appears that most investors are willing to accept the solution of investing the collateral in US treasuries despite their low returns. The situation may change when Libor levels increase again.

Increasingly, transparency is demanded by investors, and this trend points clearly to greater transparency of the structure and the composition of collateral. The situation where investors do not have timely access to information about what assets are part of the collateral is not going to be accepted for much longer.

Weather Derivatives

THE BROADER DEFINITION OF INSURANCE-LINKED SECURITIES

The term “insurance-linked securities” (ILS) has, from the very beginning, included financial instruments that technically do not contain insurance risk at all. The reason is that the risk in these securities is of the type that would often be borne by insurance companies. An ILS tied to the level of a longevity index is an example of such a security. Weather derivatives fall in this category even though they are often treated separately from other ILS as a distinct asset class. The risk of weather events leading to economic and other damage is common in insurance and is often the main risk in an insurance coverage. Some types of that risk can be dealt with more efficiently by transferring it in the form of non-insurance financial instruments such as weather derivatives. This chapter looks at weather derivative types, pricing and investing.

WEATHER DERIVATIVES DEFINED

Weather derivatives are derivative financial instruments whose payout depends on the value of a weather-related index or event. By definition, weather derivatives are not contracts of insurance. The underlyings are not financial assets with a defined price but rather variables linked to weather phenomena, such as temperature, precipitation or wind. In this sense, they are similar to catastrophe derivatives (described earlier), which do not have an underlying asset with a defined price either. Some might even put catastrophe derivatives such as wind futures in the category of weather derivatives despite their triggers being tied to insured losses from specific weather events.

The weather derivatives market appeared in the late 1990s and has grown significantly since then, even though the growth has not been steady. The market has achieved a significant degree of standardisation; the growth in exchange-traded weather derivatives, in addition to the over-the-counter instruments, has taken the market to the next stage. The types of contracts

span the same general universe as that of the traditional derivatives, with calls, puts, swaps and forwards possible.

The most common types of weather derivatives are the temperature-linked ones that provide hedging protection to utilities and other energy complex companies in case of temperature-related economic losses due to unexpectedly high or low demand for electricity, oil and natural gas. They have proved their value in this industry and have become a standard part of the tool box of instruments for managing risk and eliminating unnecessary earnings volatility.

Weather derivatives and weather insurance

There is a strong link between weather derivatives and insurance that covers damage from weather events. In fact, there is a significant overlap, since many transactions can be done in the form of either insurance or a derivative. Insurance companies wanting to be on the other side of a weather derivative transaction have sometimes even utilised a transformer to provide protection, in the form of insurance, against an adverse economic effect of weather events.

There are, however, some differences between weather derivatives and weather insurance. The main difference is that for derivatives there is no need to demonstrate actual loss. In insurance, the existence of insured loss is necessary for a claim to be paid. In many insurance contracts, there is some moral hazard involved, which is unlikely to be the case in weather derivatives. One more difference is that weather insurance is a hedge against direct losses suffered by the insured due to adverse weather-related events. Weather derivatives can also be used to hedge against indirect economic losses resulting from the weather being better than expected. This might not be common but is different from the traditional insurance approach of insuring only damaging effects of bad weather (storms, floods, other extremes). This distinction is often pointed out as a differentiator between the two types of financial instruments. However, we can usually create an insurance product that will replicate the economic effect of even this type of a weather derivative. An important additional difference is that a derivative transaction can be undone, while it is impossible to do the same in a direct fashion for an insurance contract.

Types of underlyings

The underlying asset being a weather parameter is what distinguishes weather derivatives from traditional financial derivatives. The underlyings

in weather derivatives typically have to do with temperature but can involve a number of other weather variables. Examples of underlyings include heating degree days, cooling degree days, maximum temperature, minimum temperature, average temperature, growing degree days, level of rainfall, level of snow, humidity, periods of sunshine, periods of time when wind speed exceeds a predetermined level, stream flow (all of the above calculated over various periods of time) and several others. The existence of these underlyings is what allows some to say that temperature and precipitation can now be traded as a commodity.

HEATING AND COOLING DEGREE DAYS

The most common type of weather derivatives are those tied to the number of heating degree days (HDD) and cooling degree days (CDD). They are calculated based on temperature measurements at a specified location that has measurement equipment and, almost always, historical temperature

PANEL 8.1 HDD AND CDD EXPLAINED

The calculation of HDD involves several steps. First, daily mean temperature T_i^{avg} is calculated for each day i as the average of the minimum and the maximum temperatures recorded for that day (over 24 hours)

$$T_i^{avg} = \frac{T_i^{\min} + T_i^{\max}}{2}$$

Daily HDD is the number of degrees by which the daily mean temperature deviates from a reference temperature. The reference temperature is 18°C in Europe and most other parts of the world and 65°F in the US; the difference between the two is a small fraction of a degree. A different reference temperature can be chosen in a bespoke transaction, particularly if the location is based in the tropical climate. For reference (base) temperature of 18°C, the daily HDD for day i is then

$$HDD_i = \max(0, 18 - T_i^{avg})$$

Similarly, for daily CDD we have the following definition

$$CDD_i = \max(0, T_i^{avg} - 18)$$

Then, the total number of heating or cooling degree days for a period of n days can be calculated by adding up daily values over the period

$$HDD^n = \sum_{i=1}^n HDD_i \text{ and } CDD^n = \sum_{i=1}^n CDD_i$$

observations. In almost all the cases this would be a station maintained by the government, which provides credibility to the temperature measurement results.

HDD and CDD are of most importance in the broadly defined energy sector, including utility companies, gas and oil suppliers, energy traders and others. HDD and CDD measure additional heating or cooling demand resulting from departures of the temperatures from their expected values. If a month or a season is particularly warm, resulting in extra electricity demand to power air conditioning and other equipment, a CDD derivative can hedge against the economic effects of this higher-than-expected level of electricity demand. Similarly, an HDD derivative can be a hedge against the economic effects of a higher-than-expected level of demand for electricity if the weather is colder than expected. Without providing more detailed explanations, it is worth noting that HDD and CDD derivatives are more volume hedges than price hedges, even though there is a clear relationship between price and volume.

OTHER TYPES OF WEATHER DERIVATIVES

HDD and CDD are the most common but not the only underlyings in weather derivatives. There are a number of others, even among the temperature-related ones. In general, temperature-related weather derivatives are of most use in the energy sector, tourism, retail and construction industries, in all of which earnings can be a function of temperature-related variables. Agriculture is another important example of a sector with exposure to temperature that can be partially hedged with weather derivatives.

Rainfall over a time period has an effect on the retail, agriculture, construction, tourism and other industries. Another type of precipitation, snow, can affect the same industries, and some others such as airlines and airports. Wind speed over a period of time can have an impact on the wind-generator segment of the energy industry, on agriculture and the retail sector. Sunshine hours over a period of time can have an effect on solar-energy generation, agriculture, tourism and retail and food industries.

There are many other examples. While the HDD and CDD contracts are common and have become fully standardised, for a less common underlying to hedge against a specific exposure, a custom structure usually needs to be created.

An example of a bespoke weather derivative transaction based on an uncommon trigger is the purchase of precipitation-linked coverage by the World Food Programme of the UN from AXA Re to provide immediate

funding in case of an extreme drought in Ethiopia during 2006. Since an intergovernmental organisation was involved, details of the transaction were made public.

PANEL 8.2 WORLD FOOD PROGRAMME/AXA PRECIPITATION WEATHER DERIVATIVE

The World Food Programme purchased the derivative to get immediate access to funding in case of an extreme drought, providing for emergency response to the consequent risk of famine in Ethiopia, in case the severe droughts of the previous year were to continue. The two primary reasons for engaging in this specific type of transaction were: (1) it was the most efficient use of donor funds (based on cost-benefit analysis); and (2) the necessary funds being made immediately available meant emergency response delays could be avoided.

Type: Call option.

Index: Bespoke crop water-stress index based on precipitation (rain) measured at 26 locations in Ethiopia formulaically converted into crop water-stress indexes and then aggregated over all 26 locations.

Strike: The above index being a specified level at the end of the season. The trigger was set at a level significantly below historical averages for rainfall during the agricultural season.

The trigger and the strike level were chosen based on potentially significant (catastrophic) losses to 17 million poor farmers in Ethiopia. The term from March through October corresponded to the agricultural season in Ethiopia.

The transaction was done in a pure weather derivative form even though it has been referred to as insurance in some of the press. In effect, it was insurance in the sense that it provided protection in case of damage due to severe weather conditions; the form of the transaction, however, was a weather derivative. This transaction did not result in a payout, since the drought that actually occurred was not that severe.

Another example would be a custom derivative to protect organisers of a large sporting event from the risk of cancellation due to adverse weather conditions. Such transactions have been done both in the form of a custom weather derivative contract and as a straightforward event cancellation insurance. Specific examples abound.

PAYOUT ON STANDARD OPTIONS

The payout for HDD or CDD weather derivatives is determined in exactly the same way as it would be for a financial derivative with a standard underlying asset. The existence of a cap on the payout in the case of the weather derivatives is the one point of difference to emphasise.

Figure 8.1 illustrates profit and loss on an HDD call option, with a cap per contract applied.¹ The horizontal axis is the value of the HDD index, while the vertical one is the profit and loss at expiration. The payment on the option is the difference between the actual number of HDD and the strike multiplied by the notional value, capped at a specified amount. It works the same way as a regular call option, with the exception that the horizontal axis is not an asset price but the number of heating degree days, and that a cap is applied to the payout. This option protects the hedger from economic losses due to a colder than expected winter season.

Similarly, Figure 8.2 illustrates profit or loss on an HDD put option, also capped, as is almost always the case in weather derivatives. It provides protection against economic losses due to warmer-than-expected winter season.

As in the case of traditional derivatives, combinations of puts and calls can be used to accomplish specific investment or hedging goals. A swap could be created as a back-to-back call and put combination. A weather swap can be of particular interest to an investor who has a directional view on the underlying. A collar can be used to obtain protection from extreme movements of the underlying in either direction. Figure 8.3 illustrates the payout on such a spread position. In this illustration, an out-of-the-money HDD call is purchased and an out-of-the-money HDD put is sold. The strikes of the two are different to establish a range of protection. This is a






Figure 8.2 HDD put option**Figure 8.3** HDD collar (option spread position)

zero-cost collar since the purchase of the call is financed by the sale of the put.

EXCHANGE-TRADED WEATHER DERIVATIVES

The appearance of exchange-traded weather derivatives was a very important step in the development of the overall weather derivatives market, in that it provided a degree of standardisation and liquidity. Eliminating or minimising the counterparty credit risk in weather derivatives has also been an important role played by the exchanges.

The market remains far from liquid except for a few types of contracts, but significant progress has been made. The volume has certainly expanded over the past several years, both in the exchange-traded and over-the-counter weather derivative products.

CME has been the most active established exchange in terms of trading volume and introducing new weather derivatives products. Some of the products that can be traded on the exchange are the following:

- ❑ US HDD and CDD options and futures for weekly, monthly and seasonal periods;
- ❑ Australian HDD and CDD options and futures for monthly and seasonal periods;
- ❑ European HDD and cumulative average temperature (CAT) options and futures for monthly and seasonal periods;
- ❑ Canadian HDD, CDD, and CAT options and futures for monthly and seasonal periods; and
- ❑ Asia-Pacific options and futures on temperature-related variables for monthly and seasonal periods (the CAT index calculation used for the region is different from the standard one).

In addition, options and futures on the indexes tracking frost and the amount of snowfall are also listed on the exchange.

Technically, the standard HDD, CDD and CAT contracts trade for a large number of locations in the US, Canada, Europe and Australia. In reality, trading is concentrated on a small number of locations; contracts for other locations are practically nonexistent on the exchanges and are typically done in the over-the-counter format.

While direct trading through the CME Globex platform is available only for futures, block trading can be done for both futures and options. Block trading has accounted for a big part of the volume, which is not a positive from the point of view of developing the overall liquidity in the market.

PRICING MODELS FOR WEATHER DERIVATIVES

Pricing weather derivatives presents unique challenges. Standard methods for derivative pricing do not apply. The Black–Scholes model cannot be used, for a variety of reasons, including the lack of tradable underlying assets, applicability of the standard random walk process to a variable such as temperature, path dependency, high degree of autocorrelation, and the existence of caps on payouts. HDD and CDD, the area of most activity, has been the primary field of research; two of the approaches developed are outlined below. For all weather variables the stochastic simulation approach appears to be the preferred one.

Burn analysis

This approach involves direct application of historical observations to obtain the probability distribution of the underlying index. For example, if 50 years of observations are available, they provide 50 data points that comprise an empirical probability distribution. The approach is usually applied in a

simpler fashion that involves calculating only one number, the average payout on the option, and sometimes the standard deviation around this mean value. The mean is then the fair value of the derivative; and adding profit/risk load to it gives the price.

Some fit a distribution, not necessarily normal, to the historical burn data. This approach does not change the fundamentals of the burn analysis.

Stochastic temperature models

Analysing the underlying first as opposed to directly modelling the option payout has obvious advantages. This analysis requires building a model for temperature. In the stochastic framework, it requires the development of a large number of realistic paths for temperature changes, which can be seen as scenarios. The choice of the random process for temperature to use in simulations is critical and affects results to a significant degree.

Traditional random walk models do not work well for temperature. They do not reflect two important characteristics of the temperature variable: mean reversion, and a significant degree of autocorrelation. There is an overlap between the two.

Autocorrelation models for temperature use observations in the previous days (typically from one to three days) in simulating temperature value for the following day. There are many ways to implement this approach.

The mean reversion models used for pricing weather derivatives have been borrowed from interest rate modelling and then modified. The discrete mean reverting diffusion model for temperature is a Markov transition model with gravitation to the mean value. In a simple model, the drift parameter in the random walk model is modified to include a degree of reversion to the mean temperature observed over many years. The historical mean temperature changes every day of the temperature time path. The variability of temperature, expressed as daily changes, is also time-dependent in that its statistical parameters vary depending on the point in the season. There is more than one stochastic model for temperature that takes mean reversion into account, but the fundamental approach is straightforward and does not change. Parameterisation of the models and using some of the bootstrapping techniques to augment the analysis do make a difference.

PRACTICAL CHALLENGES IN PRICING

While it is easy to develop a mathematical model for the stochastic temperature process, mundane practical difficulties complicate its implementation. Below we outline some of these difficulties.

Data issues

Historical information on temperature observations over many years for numerous weather stations around the world is a treasure trove of data to be used in modelling. It has an additional advantage of credibility, since in most cases the data comes from official government sources. However, the issues with data quality are widespread and have been universally recognised. In the US, there are thousands of weather stations that collect weather information such as temperature, precipitation, wind speed and humidity. Almost all of them have extensive records of historical data. But the reliability of the data is questionable for many of the stations. To minimise the issue of data quality, the data stations chosen to measure weather variables in weather derivatives tend to be some of the so-called first-order weather stations. This selectivity provides greater confidence in the measurements that will be the foundation for actual derivatives payouts; it also provides more reliable databases of historical observations.

Despite the choice of the more reliable weather stations, data issues remain. For example, there is a statistically significant trend in the temperature measured by many stations. It cannot be explained by climate changes or global warming because the trend often goes back several decades and is of an unusual magnitude. The explanations could vary. One of the common ones has to do with more housing being built closer to the weather stations, contributing to an increase in the ambient temperature. In addition to trends, there are also many discontinuities in the data. Some of them can be explained by a simple movement of the weather station, say, 20 years ago; a movement by 100 feet can have a noticeable effect on the measured weather variables. Data quality issues like these are so significant that simple burn analysis cannot be relied upon in many cases. Using the historical data to parameterise a more sophisticated stochastic model could introduce a significant source of error.

There have been attempts, largely successful, to detect the problems with data at individual data stations, with a focus on the first-order stations. Attempts to adjust the data to correct the discontinuities and unreasonable trends have also been made. They have added significant value but have not eliminated the data concerns in pricing weather derivatives.

Choice of time period of historical observations

While we may be tempted to use the whole historical database of observations that might go back over a hundred years, the data-quality issues identified above argue against it, or at least against giving the same weight

to every year in the historical record. Additional considerations such as the climate change or cyclical atmospheric processes affecting temperature call for making adjustments to the historical data or incorporating these current effects in another fashion. The choice is often based on judgement.

Weather forecasts

Since in weather derivatives we often deal with short time horizons, it is possible to use meteorological forecasts in pricing. When this is done, in practice these forecasts are usually used to adjust some of the model parameters, and, based on judgement, to assign weights to historical (possibly trended and otherwise adjusted) data and to the forecasts, with the weight assigned to forecasts reducing as the time horizon increases.

INVESTING IN WEATHER DERIVATIVES

There are investment funds that have the sole mandate to generate returns by trading in weather derivatives. There are funds with a broader mandate that choose to allocate some of their assets to weather derivatives. There are investors or hedgers who participate in the market in order to hedge some of their exposure to weather risk. Trading desks at some energy companies might sometimes be in the pure investor category as opposed to having their primary focus on hedging risk. The reasons for participating in the weather derivatives market differ, and investment strategies differ with them.

For investors, portfolio management presents significant challenges. The correlation among the securities in a portfolio cannot be easily ascertained. The small number of sites for HDD and CDD derivatives limits the ability to make correlation among sites a friend instead of an enemy, and to take advantage of it to manage an investment portfolio to maximise its risk-adjusted return. The lack of good tools to quantify portfolio risk exposure is a challenge to all investors in this asset class. It is encouraging that some portfolio management tools have been developed; as their sophistication and credibility increase, so will the ability of investors to analyse their portfolios and make more informed decisions.

Combining weather derivatives with other types of investments has created value for some investors. It is a natural fit with traditional ILS as well as with securities linked to emissions trading. All of these, to some degree or another, have a relatively low degree of correlation with traditional financial assets.

Specific investment strategies

Some of the investment strategies incorporating weather derivatives are the following.

- ❑ Taking advantage of mispricing of individual securities. This could be security-specific, especially if an over-the-counter and difficult-to-analyse security is involved, and the investor has the skill to analyse it better than others. It is also possible to have a directional view on a weather variable such as temperature that is different from the one implied by prices in the market. This view can be based on superior analytical tools and access to experts. It results in directional trading, the degree of which is a function of the investor's conviction level of their directional view on the market.
- ❑ Taking advantage of being able to properly capture portfolio risk. If an investor is able to reflect correlation among securities in the portfolio and quantify the portfolio exposure, they are in a position to actively manage the portfolio by buying or selling positions with the goal of increasing the risk-adjusted return of the total portfolio. Ability to actively manage a portfolio is a critical competitive advantage in the market where correlations are difficult to quantify and inefficiencies are relatively common.
- ❑ Using weather derivatives as a hedge against other positions in the portfolio (such as commodities) could be part of a broader investment strategy. In fact, it might be possible to use commodities as a hedge against some of the risk in a portfolio of weather derivatives. We should be aware that the correlation among weather derivatives and commodities has not been proven to be at the level that allows effective implementation of this strategy and profitable cross-market trading. There are obvious pair trades, however, such as a stock of a snowmobile manufacturer or a ski resort, paired with a weather derivative linked to the level of snowfall.

These are just some of the examples of investment strategies involving weather derivatives; others exist as well.

Valuation

Valuation of weather derivatives in an investment portfolio presents the type of challenge common for valuing securities that have limited liquidity. It is not always possible to ascertain the market value of a weather derivative, even for the most popular HDD/CDD types. At the same time, marking to model carries with it inherent dangers that are best avoided in security valuation.

The approach used by many in practice is to mark to market some of the most liquid HDD and CDD contracts (most popular locations) and mark to model those that are illiquid. It goes without saying that the model has to be updated to take into account new information; it cannot be based just on the original pricing analysis. For those weather derivatives that fall between these two extremes, a combination of the mark-to-market and mark-to-model approaches is often used, with greater weight typically given to the mark-to-model approach. Some employ the mark-to-model approach with a reasonability check in the form of comparing the results with available prices for similar instruments; any significant discrepancy then has to be explained, and the difference from the market indications has to be justified. The issue of valuation is also a function of the accounting rules in the relevant jurisdiction.

EMISSIONS TRADING

Emissions trading is tangentially related to weather derivatives and is not part of the ILS marketplace. The reason it is mentioned here is that it shares with ILS the low degree of correlation with traditional financial assets. As such, it can also act as a diversifier in a broader investment portfolio. In addition, it is possible to combine ILS and emissions-linked securities in the same portfolio or fund with investment returns exhibiting a low degree of correlation with global financial markets.

Emissions trading has to do primarily with commitments of parties to the Kyoto Protocol to limit or reduce their overall greenhouse gas emissions. The agreement allowed countries that have not fully used their emissions quotas to sell the excess to those who are finding it difficult to meet their targets. Typically, trading in emissions is referred to as carbon trading since carbon dioxide is the principal greenhouse gas. In addition to the emissions units (so-called AAUs), the Kyoto Protocol system allows trading of related securities such as removal units (RMUs), certified emission reduction (CER) and emission reduction units (ERUs).

There are some emissions trading systems at regional and national levels, of which the one in the EU is the largest. Recently, there has been criticism of the EU emissions trading system as ineffective in achieving the goal of significant reduction of greenhouse gas emissions, along with concerns that the proposed adoption of the same system in other markets is misguided. Under the system, companies in sectors such as energy, steel and manufacturing are given allowances for their greenhouse gas emissions, with allocations being reduced over time. If the companies do not reduce their

emissions accordingly, they are forced to buy additional permits from others to remain within their quotas. If the system is changed due to the criticism, it will only increase the volume of emissions trading and contribute to the growth of the market.

In addition to the emissions trading described above, there are also voluntary markets for emissions. In the US, which is not part of the Kyoto Protocol, one of the very first emissions trading systems was implemented to reduce SO₂ emissions. In 2009, the Regional Greenhouse Gas Initiative was implemented in nine states in the US to limit carbon dioxide emissions from power generators in the form of a cap-and-trade programme. There are several other emissions trading systems in the US and around the world.

The Chicago Climate Exchange (CCX), part of Climate Exchange plc, provides a platform for trading emissions under a voluntary system. While the system is voluntary, it is contractually binding. Parties to the agreement, which include mostly commercial enterprises but also states and municipalities, have committed to annual emissions reduction targets. Opportunities for trading arise when one party has surplus allowances while others need to buy additional ones to avoid violating the agreement. All six greenhouse gases can be traded on the exchange. The security being traded is the Carbon Financial Instrument (CFI) contract. Chicago Climate Futures Exchange (CCFE), a subsidiary of CCX, is an exchange for environmental derivatives such as futures and options on emission allowances. CCFE provides a platform for trading products ranging from futures and options on the CFI contracts to futures and options on SFIs (sulphur financial instruments).

There may be significant changes coming in the US, whether or not the country joins the Kyoto Protocol or a similar agreement. This will lead to market growth and new opportunities for investors.

TRENDS AND EXPECTATIONS

The weather derivatives market has grown very fast since its birth in the late 1990s. The overall volume and the number of trades have grown far beyond any initial expectations. The market is here to stay and will likely continue to grow as the effectiveness of weather derivative hedges becomes better understood by companies affected by weather, and as the investor community, in its search for uncorrelated assets, becomes more involved in the market.

In 2008, however, market growth stalled. In 2009, the lack of measurable growth remained the reality. The slowdown that started in 2008 is likely due only to the effects of the global economic crisis and the general slowdown in

capital markets as well as the deleveraging and reassessment of risk. The long-term prospects for the market remain very positive. This expectation is based on the following conditions and long-term trends.

- ❑ The education process will continue. A large part of the global economy is exposed to weather risk, and the ability to hedge the risk effectively with weather derivatives can give some companies a competitive advantage, leading others to follow suit. The market will expand well beyond the energy sector and related industries.
- ❑ The expansion to new industries may be a slow process, since the use of derivative products and the understanding of the concept of derivatives in general is absent in many sectors of the economy. However, long-term prospects are bright, and there is every expectation that the market will continue to expand to new industries.
- ❑ While the energy sector is already the single biggest participant in the weather derivatives market, in absolute terms its involvement is expected to increase, since there are derivative contracts that very directly address the risk of economic losses due to weather in this sector.
- ❑ The growth in the exchange-traded segment of the weather derivatives market is inevitable if the market overall is to develop. While short-term predictions in this area are particularly difficult, in the long term exchange-traded weather derivatives will likely grow more than the over-the-counter segment of the market. Liquidity, lower execution cost and overall greater efficiency are important advantages of exchange-traded weather derivatives. Another advantage, whose significance has grown, is the limited counterparty credit risk associated with exchange-traded contracts.
- ❑ Smaller companies that have rarely participated in the market will have more opportunities to do so – due not only to the general education about weather derivatives, but also to the credit risk concerns of their potential counterparties being minimised if exchange-traded products are used.
- ❑ It is possible that agriculture will be a bigger participant in the weather derivatives market if there are changes in the structure of government subsidies common in the industry. It has been suggested that the subsidies often decrease the incentive to hedge weather risk in this industry. This suggestion has not been confirmed.
- ❑ The role of insurance companies in the weather derivatives market is likely to grow again, following the retrenchment in recent years, even though insurance and reinsurance companies were some of the first

participants in the market. It is unclear what form this role will take, since it is possible that insurance companies will start using weather derivatives for hedging their insurance exposure, as opposed to always being on the risk-taking side.

- ❑ The availability of effective standard modelling tools and the access to adjusted weather data lower the barriers to entry for this market and reduce the level of information asymmetry among its participants. This is one of the most important developments conducive to growth of the market.
- ❑ Correlation between weather derivatives and commodity derivatives will be examined more closely, and this could lead to a broad use of weather derivatives in managing commodity investment portfolios. Investors can also use known correlations for pair trading of stocks of companies exposed to significant weather risk, and weather derivatives tied to that risk.
- ❑ The improved ability to model weather risk and its correlation to economic results will continue to be an important growth driver. The meteorological models for probabilistic temperature prediction will continue to improve, increasing the efficiency of weather hedges and the pricing of weather derivatives. The improvement of existing models and development of new ones for weather variables other than temperature will make it easier and more efficient to create weather derivatives based on these variables.
- ❑ The overall growth of the weather derivatives market – from increased liquidity, to adding new measurement locations, to the broader introduction of new weather variables – makes it more attractive to sophisticated investors by improving their ability to manage weather derivatives on a portfolio basis.
- ❑ The advantages mentioned above, along with the development of models able to properly capture correlation among specific weather risks in a portfolio, will give investors new opportunities to actively manage their portfolios. Low correlation of these securities with traditional investments such as stocks and bonds will further increase the attractiveness of weather derivatives as a diversifier in an investment portfolio and a way expand the efficient frontier.

1 In Figure 8.1, the option premium shown could be more precisely replaced by the premium plus the interest accrued to the time of expiration.

Part III

Securities Linked to Value-in-Force Monetisation and Funding Regulatory Reserves

Funding Excess Insurance Reserves

EXCESS INSURANCE RESERVES

Regulations governing the way insurance liabilities are calculated can differ, sometimes materially, from the standard GAAP or IFRS principles. Economic reserves, defined as those based on best estimates, could differ from both the statutory and GAAP liabilities. This is particularly true in the US, where statutory accounting rules could add a significant degree of conservatism to the level of insurance reserves. It is hard to find examples of such significant divergence outside of the insurance industry. Economic reserves, defined as those based on best estimates, could differ from both the statutory and GAAP liabilities.

Prudent regulation can sometimes result in balance-sheet liabilities substantially in excess of economic reserves. This situation can create capital strain on insurance companies that, from their point of view, is not justified and is not supported by the economic theory. Reserves that are higher than economically necessary decrease the probability of insurance company insolvency in the short term. At the same time, they put downward pressure on shareholder returns and can force insurance companies to raise rates on the products with high reserve requirements.

Contrary to a common belief, the problem of having to fund reserves that appear excessive from the economic point of view is not limited to life insurance. It can arise in both life and property-casualty insurance, as well as in the annuity business. This chapter shows how insurance companies can use capital markets solutions to bring about surplus relief and reduce the cost of capital. The chapter presents examples of funding solutions for specific products, which demonstrate potential structures that can be used in reserve funding in other situations.

SOME EXAMPLES

Five examples of insurance products that in some cases have required the establishment of reserves considered to be excessive are presented below.

- ❑ Level-premium term life insurance contracts in the US are subject to Regulation XXX, which was adopted by the National Association of Insurance Commissioners (NAIC) in 2000 and became effective in most states in 2001. Regulation XXX requires the use of a mortality table that is considered to be overly conservative even by the rating agencies, leading to statutory reserves being far in excess of economic reserves and creating a capital strain on a company engaged in writing this type of life insurance.
- ❑ Establishing liabilities for universal life insurance policies in the US is subject to Actuarial Guideline 38, also known as AXXX. It too is believed to have imposed overly conservative standards on the reserve calculations, setting up liabilities at levels far exceeding economic reserves needed to fund company obligations under the insurance contracts. Universal life insurance policies with secondary guarantees are the ones negatively affected by these requirements. While some regulatory changes are having the effect of reducing the overall level of AXXX reserves, there remains a sizable gap between statutory and economic reserves for such policies.
- ❑ Motor insurance in Europe under the current regulatory regime, which will change and is already changing, is subject to accounting rules perceived by some to require, in certain cases, the establishment of reserves in excess of those economically necessary. This will possibly remain the case until Solvency II is fully implemented, and maybe even after that.
- ❑ Long-tail lines of casualty insurance are in most cases supposed to be reserved based not on the present value of expected future loss payments but rather at full value, without taking into account time value of money, in jurisdictions such as the US. Certain losses related to lines of business such as workers' compensation insurance can be discounted, but there are many situations where statutory regulations do not allow loss reserve discounting; this limitation could result in a significant difference between statutory and economic reserves.
- ❑ Variable annuity contracts with secondary guarantees can necessitate establishing reserves at the levels deemed excessive, also resulting in capital strain and reduction of statutory surplus. The move to principles-based reserving can alleviate some of this strain.

Other examples can be brought up as well. It is important to note that all of them are jurisdiction-specific.

“EXCESS” RESERVES

“Excess” reserves may be seen as unnecessary by the proponents of using economic values for every item on the balance sheet, but they can serve an important role from the point of view of regulators and policyholders. Regulators may not see these reserves as anything that can be called “excess”, but rather could regard them as essential in maintaining insurance company solvency and protecting the interests of policyholders.

Insurance companies often argue that the stringent reserving requirements that create the kind of capital strain described above actually hurt policyholder interests. Keeping additional capital requires charging higher rates to maintain the same return on capital. Consequently, consumers suffer. This logical argument has an equally logical counterargument that consumers suffer when insurer insolvencies happen, either directly or in the form of having to pay higher insurance premiums to cover state guarantee fund assessments. There is a continuing disagreement as to where the right balance between the two should be struck.

There is also a disagreement as to whether the reserves are truly excessive even on an economic basis, and, if they are, to what degree. This is a particular issue in reserving for long-tail casualty insurance lines of business, where some companies engage in implicit discounting by understating the value of liabilities on their balance sheets. In such cases, balance-sheet reserves might not be overstated from the economic point of view.

In cases where reserve discounting is allowed and even mandated, there could be a disagreement over the right discount rate to use. Excess of statutory balance-sheet reserves over economic reserves can exist if statutory accounting rules specify a discount rate lower than what might seem reasonable from the economic point of view. The choice of proper discount rate is subject to judgement, possibly resulting in a disagreement about whether the reserves are excessive or not.

FUNDING SOLUTIONS

An insurance company finding itself required to establish “redundant” reserves might attempt regulatory arbitrage by finding a way to transfer liabilities to a jurisdiction with less demanding accounting rules without violating its own domicile regulations. Doing so, however, is not always possible. Another way to fund the “excess” reserves is through securitisation or a lending arrangement. This too is effectively a form of regulatory arbitrage, and, though performed in a somewhat different fashion, it similarly could involve the transfer of liabilities to a different jurisdiction as part

of the securitisation structure. A bank lending structure with the use of reinsurance can be another solution in some circumstances.

A company that has discovered a way to relieve capital strain caused by the requirement to hold “excess” reserves finds itself in the position of competitive advantage. This competitive advantage is unlikely to be sustainable if the funding solution is easily available to other companies with the same product lines. However, any company left behind and unable to fund its “excess” reserves in a cheaper way will certainly be put in an unfavourable competitive position.

A funding solution such as liability-related securitisation can provide an efficient way to alleviate the capital strain in a cost-effective fashion. Securitisation or a bank lending arrangement, whether or not it directly involves the use of a reinsurance mechanism, is an important tool to consider when faced with such a capital strain issue.

Often, this tool can be used for more than one reason at the same time. For example, it can involve reasons including true risk transfer – such as when the reserves might be excessive on the expected net-present-value basis from the economic point of view, but there is also significant volatility around the expected value, and the insurance company wishes to transfer this risk to a third party such as capital markets investors.

EMBEDDED VALUE AND VALUE-IN-FORCE SECURITISATION

Embedded-value securitisation or monetisation, described in other chapters, could be seen as another example of funding liabilities that are set up on the balance sheet by accounting rules and arguably do not reflect economic reality since they do not follow the rule of matching the time of expense and revenue recognition.

For example, most of the expenses involved in originating life insurance policies are front-loaded, and in the beginning the cashflows to the insurance company are negative. The GAAP concept of deferred acquisition cost (DAC) is not always recognised by statutory accounting rules; consequently, an insurance company might have to immediately fund the cost of originating the policies. The fact that profits are expected to emerge later from such insurance policies does not negate the requirement of immediate expense recognition. This requirement creates capital strain associated with writing new business; the better a company is doing in marketing and selling its products, the worse its statutory financial results might look. To provide surplus relief, the company might sell to investors some of the future cashflows from the policies on its books in return for immediately

receiving cash instead of having to wait for the emergence of profits from the policies.

A fast-growing company can find itself under significant capital strain; securitisation or a bank funding arrangement is a capital management tool, as much as or even more than traditional reinsurance, that serves the purpose of providing surplus relief. This type of securitisation is referred to as securitisation of embedded value (EV) or value-in-force (VIF). It is an efficient way of accelerating the balance sheet and relieving the capital strain caused by writing new business. Many of the EV or VIF monetisations have been performed in the M&A context, generating cash needed to finance the acquisition, or in the context of demutualisation.

This topic is treated in more detail in Chapter 10, which describes specific structures used in securitising or monetising embedded value, as well as providing an illustration of an embedded value securitisation. Some of the accounting considerations are also described there.

MARKET FLUIDITY

Regulations leading to the establishment of “redundant” reserves change as the whole regulatory framework continues to evolve. Moving to principles-based reserving is likely to significantly reduce the level of reserve redundancy. Most of such regulatory developments are now originating from Europe, but the US regulators are also working on modernising the existing regulations. This is expected to be a difficult multi-year process. The landscape is constantly changing. Some funding solutions are no longer feasible in the current environment. For example, financial guarantee, an essential part of some securitisation solutions to funding excess reserves, is no longer available and is unlikely to become available for several years at a reasonable cost.

As regulations change and the financial environment changes as well, some excess reserve funding challenges disappear, either for good or only to appear later in a different form, and new challenges sometimes surface. Funding solutions will continue to develop, too, either because the old ones no longer work or because there is a need for such solutions in a new area.

RBC REQUIREMENTS LEADING TO “UNNECESSARY” CAPITAL STRAIN

Capital strain due to establishing reserves considered by many insurance companies to include an excessive degree of conservatism is also shown in the way risk-based capital (RBC) is calculated according to the National Association of Insurance Commissioners rules in the US. Variable annuities

with secondary guarantees are one such case, showing that the need to maintain sufficiently high RBC levels could create capital strain on an insurance company similar to how it could be created in a more direct way by increasing the level of statutory balance-sheet reserves.

The principles-based reserving approach can help to alleviate the RBC strain as well. It provides the degree of realism typically not achievable in any formula-based approach. By its very definition, principles-based reserving utilises risk analysis methods to quantify risks, including stochastic modelling where necessary; captures all the relevant risk factors, including guarantees embedded in insurance or annuity contracts; and allows more extensive use of company-specific assumptions where appropriate. This approach reduces the chances of reserves being “redundant”. The NAIC principles-based valuation project is expected to bring the regulation closer to adapting some elements of principles-based reserving. The RBC C-3 Phase II has established an important precedent in the US of the use of principles-based methodology. The proposed RBC C-3 Phase III will bring the industry another step closer to principles-based reserving.


The approach being adopted in Europe is much closer to true principles-based reserving, but it still leaves a lot to be desired. In addition, in Solvency II, some of the criticism of “market consistency” has merit, as well as the criticism of the possibly excessive emphasis on the one-year horizon.

Capital markets solutions can be used to alleviate the “RBC strain” and improve the level of its risk-based capital. Securitisation is one such solution.

REGULATION XXX RESERVE FUNDING

The Valuation of Life Policies Regulation in the US, also known as Regulation Triple-X or XXX, established statutory valuation requirements for most life insurance products. Its effect was felt most in calculating reserves for guaranteed level-premium term life insurance policies. Regulation XXX has resulted in insurance companies’ having to increase, by a sizable amount, the level of reserves they set up for new level-premium term policies.

The adoption of Regulation XXX created a gap between statutory reserves and economic reserves. Economic reserves are based on best estimates and do not have the safety margin that is included in statutory reserves to ensure that future policy obligations are met. The gap or “redundancy”, illustrated in Figure 9.1, could be a multiple of the statutory reserves, especially when several years of production are considered. Premiums can be guaranteed for as long as 30 years, leading to the need to have higher capital against these reserves for a long period of time.

Figure 9.1 “Redundant” reserves created by Regulation XXX

Non-US reinsurance companies are not subject to the Regulation XXX requirements. A natural choice for a US life insurance company would be to engage in regulatory arbitrage by reinsuring some of its level-term life insurance book to a non-US reinsurer. However, to receive reinsurance credit and resulting reserve and capital relief, reinsurance has to be fully collateralised by qualifying assets held in a trust. An alternative is a letter of credit (LoC) provided by a bank or another financial institution. The use of letters of credit has been common in providing reinsurance collateral, especially before the credit crisis that started in 2007. A typical letter of credit has the term of one year. Longer-term letters of credit are available but carry a higher cost. Letters of credit with a term of 20 or 30 years are extremely uncommon and expensive. This is the term for which reinsurance should be in force to provide the necessary reserve relief. Long-term letters of credit, even when available, are so expensive that paying for them does little to alleviate the capital strain. The option of obtaining reinsurance for a short term such as one year, with the intent of then renewing the reinsurance contract on an annual basis, carries with it the risk that in the future reinsurance might not be available, at least not at the anticipated cost. This risk is taken into account by rating agencies and investors in their analysis of insurance companies. Regulators are also aware of the risk. Not having a longer-term solution, and relying on short-term reinsurance and short-term letters of credit, effectively creates financial leverage for insurance companies.

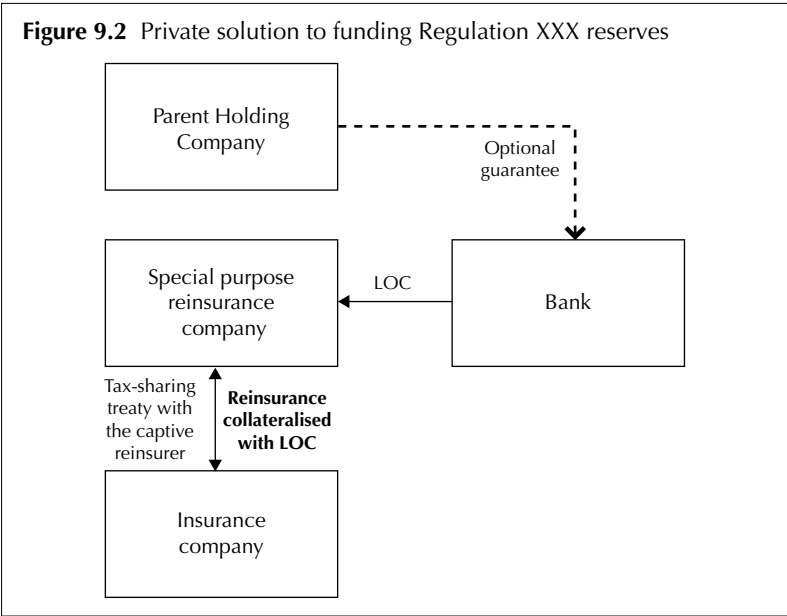
A funding solution for Regulation XXX reserves would address all of these problems and reduce the capital strain on writers of guaranteed level-premium term life insurance.

LETTER-OF-CREDIT FACILITY FOR FUNDING REGULATION XXX RESERVES

A letter-of-credit banking facility has been utilised to fund, in a relatively inexpensive way, the “redundant” reserves created by Regulation XXX. Letters of credit, if their term is sufficiently long, are usually treated by rating agencies as operating as opposed to financial leverage. One of the advantages of this way of funding redundant reserves is the lower execution cost when compared with the securitisation solutions described later. Figure 9.2 illustrates how such a credit facility can be structured.

A special purpose reinsurance company is formed as a captive of the life insurance company seeking reserve relief. The reinsurance company enters into a co-insurance agreement with the primary insurance company. The reinsurance collateral is the letter of credit from a bank. The agreements provide for automatic extension if certain conditions are met. They can also allow for the arrangement to be terminated if the statutory regulations change and no longer require the maintenance of excess reserves, or if the tax code changes and the structure becomes less tax-efficient. Typically, there would be a tax-sharing treaty between the operating insurance company and the captive reinsurer.

The structure can include an optional guarantee from the parent company to reimburse the bank if reserves do end up being deficient and the letter of



credit is drawn. Rating agencies have a negative view of the parent company's providing such a guarantee (as they would have of most recourse arrangements). Ways to avoid the need for a parental guarantee from the holding company are to demonstrate through modelling that the risk is minimal and to contribute additional capital to the captive, either directly or through a specific clause in the reinsurance treaty. The latter solution, however, makes it more expensive to the insurer and could offset some of the benefits of the funding structure.

The private nature of the arrangement has some benefits not available in a traditional securitisation described below. There is the flexibility in the terms and conditions that allows a greater degree of customisation. There is never a need for a third party to provide a financial guarantee, assuming the bank issuing the letter of credit has sufficiently high ratings. Cashflow modelling and other actuarial analyses are not as extensive as in the case of securitisation. Finally, there is no need to obtain a rating on the notes from a rating agency.

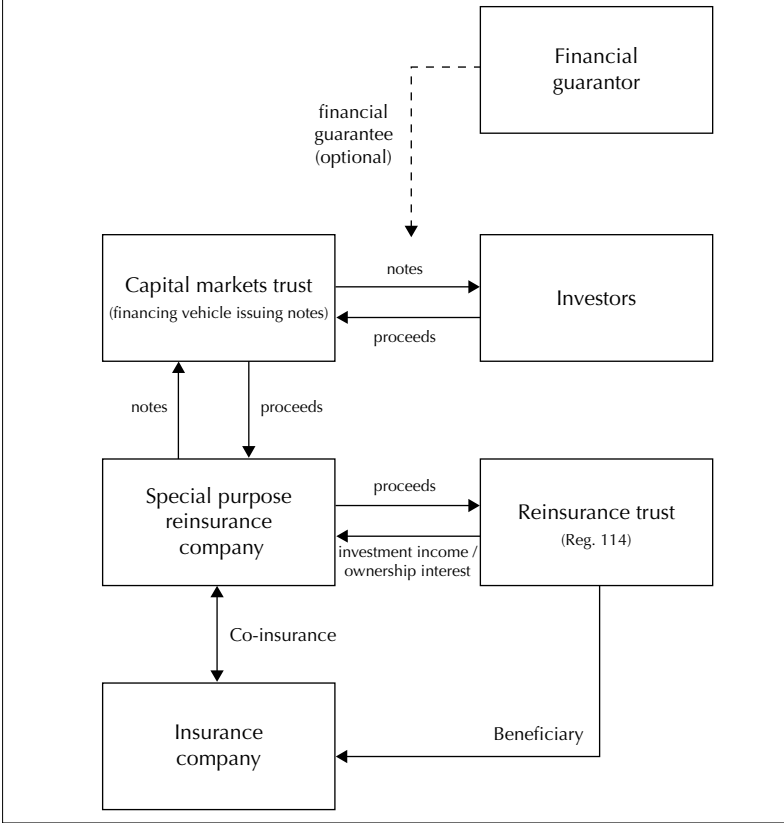
Banks that accumulate this type of risk on their balance sheets might face a problem if they are not able to pass the risk along to investors, either directly or in repackaged form. There are banks that currently hold billions in XXX risk.

SECURITISATION OF REGULATION XXX RESERVES

One solution to the XXX challenge is securitisation, which has been employed a number of times. Regardless of whether we are going to see more of such XXX securitisations in the future, reviewing the structure of such a securitisation is instructive in understanding the ways of funding redundant reserves.

A representative basic structure of a securitisation of the excess reserves is presented in Figure 9.3. A special purpose reinsurance company is established by the operating insurance company as its captive, or, in some cases, this could be done by the holding company. Several transactions are then entered into simultaneously. Pursuant to a reinsurance agreement, a significant portion of the excess reserves for guaranteed level-premium life insurance policies is then ceded to the reinsurer. A finance vehicle, an SPV, is formed to issue securities to investors to fund the part of the excess reserves ceded by the operating insurance company to the reinsurer. This capital markets trust passes on the proceeds to the special purpose reinsurance company in exchange for surplus notes of the reinsurer. These proceeds allow the reinsurer to establish collateral in the form of qualified

Figure 9.3 Securitisation structure for funding Regulation XXX reserves



assets deposited in a Regulation 114 trust. There are certain requirements as to the quality of assets in the trust; these are intended to lower the credit risk. The assets in the trust are held solely for the benefit of the ceding company, while the reinsurer is the grantor of the trust. Regulation 114 also requires that the trustee be a qualified financial institution.

Payments to investors are funded by the dividend payments from the special-purpose reinsurance company to the issuer. The reinsurer is able to pay these dividends drawing from the cashflows received from the reinsurance trust as the excess reserves are released; from the investment income on the assets in the trust; and from the cashflows from the ceding company paying premiums under the reinsurance agreement.

The securities are non-recourse, differentiating them from the private solutions where at best the securities have limited recourse.

The reinsurance contract should involve sufficient degree of risk transfer for the contract to be afforded reinsurance accounting treatment. Deposit accounting would negate the benefits of the structure.

In the past, financial guarantee was used to enhance the ratings of the notes issued to investors. AAA rating on the tranches covered by a financial guarantee made the securities attractive to a wide universe of investors that invested based more on the financial guarantee than on their having performed any analysis of the insurance risk. It is unlikely that such financial guarantee will be available in the future, at least at the cost that makes sense to all parties.

Other ways to increase the ratings of the securities, or to provide a greater level of comfort to investors when the notes are not rated, are to perform more rigorous actuarial analysis and to put additional equity in the special purpose reinsurer. The level of overcollateralisation plays an important role in the analysis.

OTHER SOLUTIONS

An example of another approach, used more often now when the financial guarantors are no longer willing to provide protection at a reasonable cost, is to obtain a financial guarantee from the holding company. The parent company would then agree to reimburse investors in case the cashflows from the reinsurer to the issuer are insufficient to cover the obligations to the investors. The higher the rating of the holding company, the greater the value of this guarantee.

There are negatives to the sponsor in utilising this solution, since rating agencies are not likely to view the guarantee favourably in assessing the financial strength of both the holding company and the operating insurance company. Depending on the details of the structure, the holding company guarantee might be viewed as being not too different from a guarantee provided directly to the operating insurance company.

ADDITIONAL CONSIDERATIONS FOR INVESTORS

Unwrapped securities can still receive an investment grade rating, albeit not AAA. The need to perform more rigorous actuarial analysis is greater for lower-rated or unrated tranches; the investor needs to better understand the risk-and-return profile to make informed investment decisions.

The transfer of the excess reserve liabilities to investors creates securities with a long tenor, between 15 and 30 years. Not all investors are interested in securities with such a long tenor. Limited liquidity adds to the risk and calls for an extra return to compensate the investor for assuming the risk.

Once the legal structure of a non-recourse securitisation solution has been vetted, investors have to assess the probability of the actual mortality and lapse experience being so different from the assumptions that the cashflows are insufficient to make payments on the investor notes. The possibility of an asset meltdown in the reinsurance trust, despite the restrictions on the asset quality and the required overcollateralisation, should also be considered. Tranche subordination is important to investors, as structurers are well aware. Without financial guarantee, the analysis requires that investors be familiar with the risk and be able to adequately assess it. In practice, the analysis is likely limited to the review of actuarial studies and the cashflow modelling already performed; but it still requires a certain degree of expertise, thus automatically excluding most potential investors. In the past, when financial guarantee was readily available, the universe of potential investors was much greater, but only for the wrapped tranches.

Stress testing and scenario testing are key to the investor analysis. Designing appropriate scenarios and assigning probabilities to each of them largely determines the risk premium that investors would charge for the securities. The key risks – mortality, lapsation, timing, investment, legal, expense level and others – have to be carefully analysed and stressed, taking into account correlation among them.

Given the limited size of the secondary market, it might also be prudent to assume that the securities will be held to maturity.

Investors should ensure that they are protected against an arbitrary action by the operating insurance company. The legal structure should provide such assurance. As in any securitisation, it is important to confirm that cashflows are stable; sensitivity analysis should provide such confirmation. Stochastic modelling, when done properly and based on reasonable assumptions, is the best way to analyse these securities. Stress testing, with specific stress tests developed, is an essential part of the analysis.

FUNDING AXXX RESERVES

Actuarial Guideline 38, also known as AXXX, was enacted to set rules for determining reserves for universal life insurance policies in the US. It is believed by many to have imposed overly conservative standards on the reserve calculations, leading to the statutory balance-sheet liabilities being established at levels far exceeding economic reserves needed to fund company obligations under the contracts. Universal life insurance policies with secondary guarantees are the ones affected by these requirements. While some regulatory changes are having the effect of reducing the overall

level of AXXX reserves, there remains a sizable gap between statutory and economic reserves for such policies.

Similar to financing XXX reserves, funding solutions have included bank credit facilities and securitisation. Securitisation, however, is more difficult for AXXX reserves than for XXX reserves. The uncertainty related to mortality, lapsation, investment and other assumptions is greater for universal life insurance reserves. Additional challenges relate to an even longer time period over which reserves run off, and possible greater correlation among the assumptions.

LOSS PORTFOLIO TRANSFER

Loss portfolio transfer could be a way for an insurance company to fund the difference between the statutory balance sheet reserves and economic reserves that result from statutory accounting rules not permitting discounting of future cashflows when calculating reserves. The difference between reserves calculated in these two ways is particularly pronounced for long-tail casualty lines of insurance. Even when discounting is permitted, the prescribed discount rate is sometimes considered to be too low.

The transfer of the reserves to an entity that can legally discount them is a potential funding solution. Such a transfer would typically be done in the form of reinsurance.

Depending on jurisdiction, however, loss portfolio transfer that achieves these economic benefits might not be allowed at all; regulatory rules differ widely in this respect. If the reinsurance company has to post full collateral in such a transaction, most of the economic benefits disappear.

Some jurisdictions do not permit discounting for specific lines of business, because of doubts in the minds of regulators as to whether the reserves and future expenditures associated with claim payments and loss adjustment expenses are adequate. If they are inadequate to begin with, discounting can lead to severe underreserving. The concern might be justified: the prohibition on reserve discounting has led some companies to understate their liabilities. This practice is referred to as implicit discounting.

Loss portfolio transfers could include significant insurance risk of adverse reserve development, which has to be carefully considered by the party assuming the liability from the insurer. Transactions intended to be finite reinsurance have all too often turned out to transfer considerable risk that was not accounted for in the analysis and pricing. Long-tail lines of business present reserving challenges and involve a significant degree of uncertainty.

CONCLUSION

Accounting rules keep changing and remain inconsistent among jurisdictions. The degree of conservatism in establishing balance-sheet liabilities demanded by insurance regulators varies from product to product. It is likely that there always will be insurance products with a noticeable gap between statutory balance-sheet reserves and best-estimate reserves, leading to a strain on surplus that might be alleviated through securitisation or some other funding mechanism.

The examples of funding solutions for specific products presented in this chapter illustrate potential structures that can be utilised in reserve funding in other situations. They show how insurance companies can use capital markets solutions to accomplish surplus relief and reduce cost of capital.

Developments such as Solvency II and the move to principles-based reserving will eventually reduce the gap between statutory and economic reserves for insurance companies. However, they will not completely eliminate the gap.

Investors in general are developing greater expertise in analysing insurance risk and insurance-linked securities. As this process continues, it could make it easier to transfer to capital markets all types of insurance assets and liabilities, including excess reserves as well.

Embedded Value Securitisation

RATIONALE FOR EMBEDDED VALUE SECURITISATION

Embedded value (EV) securitisation is the exchange by an insurance company of its future profit stream on an existing book of insurance business for a monetary consideration received from investors now.

The idea of “accelerating” profits is not unique to the insurance industry. In fact, the concept is more often used in other industries than in insurance. There could be a number of reasons for an insurance company to enter into such a transaction. A securitisation of future cashflows from a block of insurance business or a whole insurance company serves the general goal of monetising the EV of the business, or at least capitalising the prepaid acquisition costs associated with writing insurance policies or annuity contracts already on the books. The insurance company receives immediate access to the value of the future profits embedded in its existing business. EV securitisations could be performed in the mergers and acquisitions (M&A) context to help fund an acquisition. When an EV securitisation is performed to alleviate the effect on an insurance company of the expense of writing new insurance policies or annuities, the advantage is twofold. It can solve the liquidity problems and reduce the capital strain caused by statutory accounting requirements of immediately recognising prepaid acquisition expenses while not allowing any recognition of expected profits until their actual emergence.

EV securitisations are often referred to as value-in-force (VIF) securitisations, reflecting the fact that the securitised future cashflows are associated with policies already in force on the day of the securitisation.

EV securitisation should present a good example of disintermediation and of insurance companies transferring the risk to investors instead of continuing to serve as giant risk warehouses. In practice, however, these transactions are not common, and the risk transferred to capital markets in such transactions is limited.

Investors are usually not willing to take on all of the VIF risks, at least not at a price that would make the transaction attractive to the insurance company. The risks that affect the emergence of profits from a block of life

insurance policies or annuities – that is, the risks that actual profits will be lower than suggested by actuarial projections – include such risks as that of the difference between projected and realised mortality rates, lapse rates and investment returns.

Embedded value securitisation or monetisation is a capital-management tool that can also be beneficial in the context of demutualisation. This chapter examines embedded value securities, describes the general process of securitising and monetising embedded value or value-in-force, analyses the reasons for entering into these transactions and provides examples of how they can be structured.

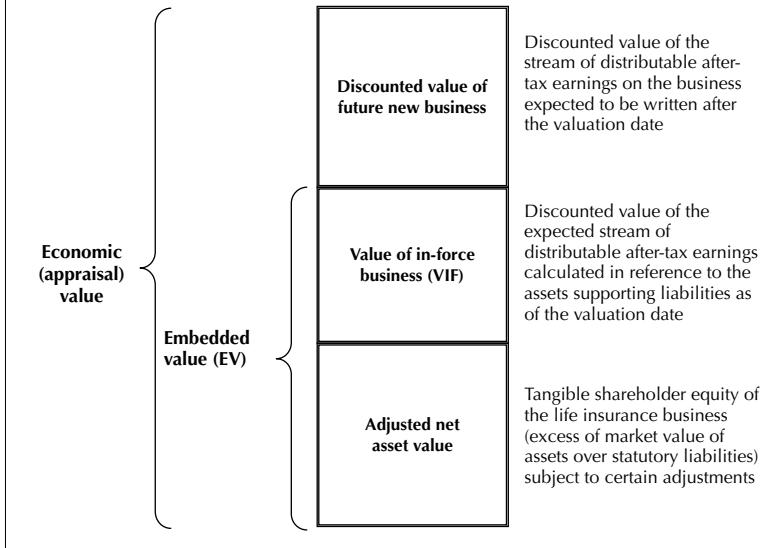
EMBEDDED VALUE AND VALUE-IN-FORCE DEFINED

There is a significant inconsistency in the way that the terms “embedded value” and “value-in-force” are defined. These concern the general questions of definition as well as the specific ways and assumptions used for calculating EV and VIF. In most cases, the terms “embedded value” and “value-in-force” are used interchangeably. For the purposes of this discussion, we will define EV as the total economic value of a life insurance business reduced by the value of the future new business.

We define adjusted net worth as the shareholders’ free surplus at the after-tax market value and the statutory capital subject to a number of adjustments. Depending on the regulatory regime, items such as asset valuation reserve, some or all of the unauthorised reinsurance and certain non-admitted assets are added, while items such as debt and surplus notes are subtracted. An important adjustment, mostly related to VIF, is the cost of capital. In this context we define cost of capital as the opportunity cost of the target surplus level that reflects the difference between the assumed future after-tax return on the surplus and the rate used for discounting this income and future releases of the target surplus.

VIF, not reflecting the cost of capital, is the net present value of the stream of distributable after-tax earnings generated by the business in-force and calculated in reference to the assets supporting the liabilities as of the valuation date.

If all the calculations were performed on a fully economic basis with the immediate recognition of expected profits, items such as VIF would not exist. The primary reason for VIF is the regulatory requirement of immediately setting up prudent reserves for life insurance policies, while the policy acquisition costs are incurred around the time of policy inception, and insurance premiums are typically paid uniformly over the term of the policies.

Figure 10.1 Embedded value and value-in-force defined

The resulting mismatch gives rise to VIF. Statutory accounting rules lead to “front-loading” capital requirements, while profits are generally “back-loaded”. Writing life insurance policies, no matter how profitable, could initially lead to a loss on a statutory basis and to capital strain on the company. The profits are recognised only over time as premiums are paid and statutory reserves are released.

EV accounting is growing in recognition, even though there is no full agreement across countries and companies on how EV should be calculated. European countries, and in particular the UK, are at the forefront of these developments.

While the discussion above and Figure 10.1 differentiate between EV and VIF, in practice the two terms are often used interchangeably.

Regardless of the specific technical details of calculating EV or VIF, it stands to reason that securitisation would help a life insurance company to get immediate access to the “hidden profits” expected to emerge in the future from an in-force block of life insurance policies, as well as to reduce the leverage created by the capital strain.

DIRECT MONETISATION VERSUS TRUE SECURITISATION

Full securitisation, as defined in previous chapters, requires true sale. In securitising EV, this means, among other things, transferring the assets to a

special purpose bankruptcy-remote vehicle. In direct monetisation, on the other hand, the assets are segregated but remain with the insurance company and could be used to satisfy other obligations of the insurer, particularly in case of insolvency. The segregated policies, including associated assets and liabilities, are referred to as a “closed block”. Typically, a closed block would not be transferred to a special purpose vehicle (SPV) and will remain with the insurer; consequently, the concept of the closed block is more commonly used in direct monetisation as opposed to true securitisation. In practice, the term “securitisation” is usually used to describe both true securitisations and direct monetisation of future cashflows from a defined block of policies. In addition, true sale in the legal sense is usually precluded by regulatory constraints.

CLOSED BLOCK

In the context of EV securitisation or monetisation, a closed block is defined as a segregated segment of the portfolio of insurance policies, typically participating or dividend paying, along with associated assets and liabilities. Only policies already on the books on the date of establishing the closed block are included. No new policies may be added to the closed block, hence the use of the term “closed”. The only possible exception is the new policies generated through the use of conversion features of the policies already in the closed block. Effectively, these policies are put in run-off and managed separately. In most cases, closed blocks have been established in the process of demutualisation.

The way a closed block is formed and managed, in particular in the context of demutualisation, is largely determined by regulatory constraints designed to protect the interests of policyholders whose policies are placed in the closed block. In addition, the way a closed block is managed is supposed to assure equitable treatment of all policies comprising the closed block. It is intended to avoid situations where the last remaining policies in the closed block receive a windfall at the expense of the policies that have expired or exited the closed block for other reasons earlier. The opposite situation, that of insufficient assets left for the last remaining policies, should also be avoided. The treatment should be equitable and consistent throughout the life of a closed block. Separate administration of a closed block is intended to accomplish this goal.

INVESTOR PERSPECTIVE

True securitisation offers obvious advantages to investors by minimising the downside stemming from insurance company insolvency or serious finan-

cial difficulties that might lead to the leakage of closed block assets. In this sense, securitisation of future cashflows is not different from most other types of insurance risk securitisation.

Overcollateralisation serves an important role in protecting the interests of investors. It could be accomplished, for example, by securitising only a certain percentage (such as 50% to 70%) of the overall expected future cashflows.¹ Other ways to protect investor interests are structure-specific and are discussed next.

SPECIFIC STRUCTURES

A number of structures have been developed for monetising future cashflows from insurance policies. Figure 10.2 provides an illustration of one such structure. In many of the completed EV securitisations, a monoline financial-guarantee company provided a credit wrap to increase the rating of the notes sold to investors. Given the general retrenching of financial guarantors and the increased cost of the protection they provide, it is likely that few, if any, EV securitisations will have such a credit wrap employed in the near future.

In Figure 10.2, a special intermediate holding company is formed between the parent holding company and the operating life insurance company. The business is split into closed-block and open-block segments, with the closed block managed separately. The special purpose intermediate holding company issues debt, usually in tranches. Different tranches have different risk and return profiles and may appeal to different categories of investor, particularly in cases where some type of credit-enhancement mechanism is employed to boost ratings of one or two tranches. Cashflows from the closed block are used to pay the interest on the debt and repay the principal. DSCA, the debt service coverage account shown in Figure 10.2, is funded at a certain percentage of the total debt from the very beginning, and is invested in high-grade corporate bonds or even government securities. The current general emphasis on minimising credit risk leads investors to seek greater levels of overcollateralisation and higher quality of securities in collateral accounts. In the past, a DSCA was funded by 20–25% of the proceeds of the securities issued to provide sufficient collateral. In addition, the collateral system typically includes security interest in the life insurance company that contains the closed block, to further protect investor interests.

Extra cashflows could be paid to the main holding company as dividends in this limited-recourse structure. The parent holding company in the illustrative structure could have other subsidiaries as well. A sizable part of the

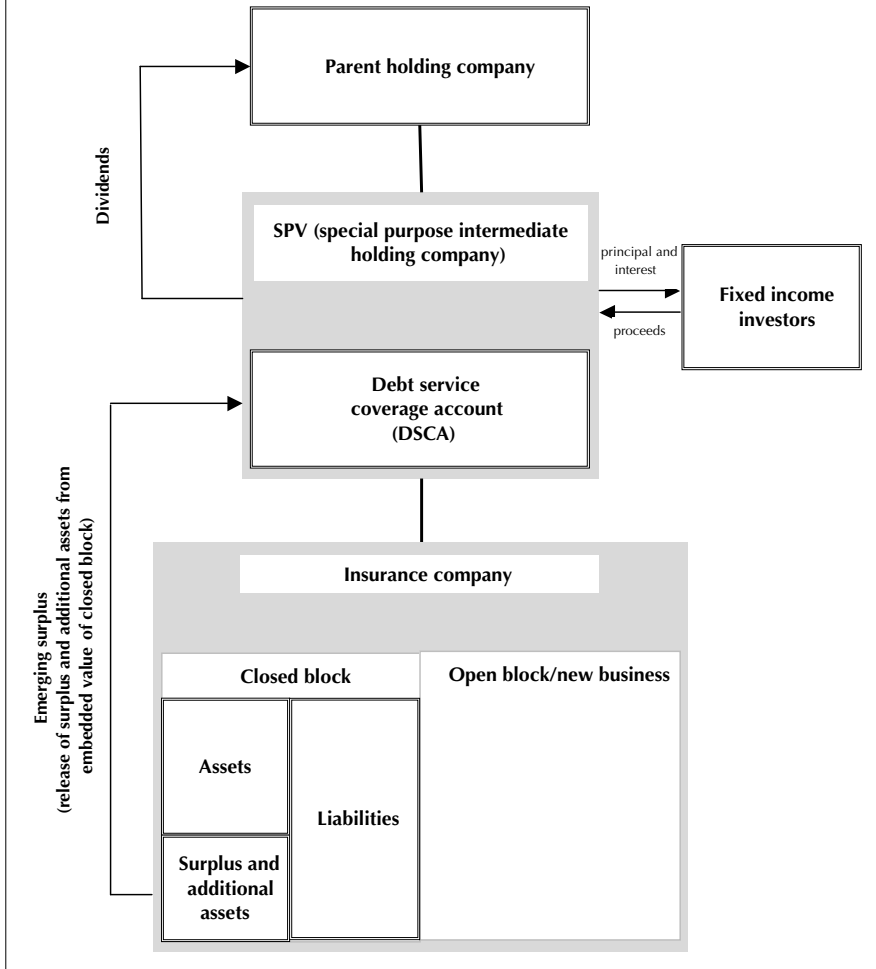
debt issued could be paid to the parent holding company. This is one of the key attractions of securitising or monetising the value locked in the closed block of insurance policies.

There are significant contractual restrictions on the activities of the operating life insurance company – in particular regarding management of the closed block – as well as on the activities of the special intermediate holding company. These restrictions are designed to minimise risk to the fixed income investors. However, risks to investors remain and could be substantial.

If designed properly, a closed block can be a source of cashflows that are relatively stable and predictable. This relative stability and predictability are the reason why such cashflows could be securitised.

Investors have several layers of protection in this and similar structures in addition to those already mentioned above. The special purpose intermediate holding company's obligations to the investors are senior to any other obligations it might have. Typically, the special intermediate holding company is not allowed to issue any other debt, even if the debt would be junior to the fixed income securities shown in this structure. A number of events, such as a downgrade of the insurance company below a certain level, could trigger the availability to debt service of additional funds – for example, those that might have been placed in a separate trust account and that represent the excess of the dividends paid by the operating insurance company over the scheduled interest payments. The covenants would generally include additional provisions to protect investor interests. The structure is supposed to remain fixed in the sense that the operating insurance company is not allowed to transfer or pledge any assets related to the closed block. Specific investment guidelines are established and should be followed as long as the investors have not been paid back. The operating company is not allowed to significantly change the nature of the business it is engaged in. The special purpose intermediate holding company is not allowed to transfer the ownership of the operating insurance company and should remain its sole owner for as long as the investors have not been paid back. There are also specific covenants intended to minimise the risk of insolvency, but such protection is better accomplished in the modified structure presented in Figure 10.3.

In general, a risk exists that the operating insurance company will encounter difficulties related to its block of ongoing business, restricting its ability to pay upstream dividends and jeopardising payments to investors. Insurance companies are heavily regulated and might not be allowed to pay dividends under certain conditions. In extreme cases the insurance company

Figure 10.2 Embedded value securitisation/monetisation structure

can be liquidated or put in rehabilitation; given the level of discretion available to regulators in the insurance industry, this can happen even if the company is not technically insolvent. Regulators could also get involved in the decisions concerning the management of the closed block to assure that the interests of the policyholders are protected and to prevent the closed block from ending up having to subsidise the open block of ongoing business. All of the above have the potential to affect investor interests.

The use of a special purpose reinsurer can alleviate some of the investor concerns. In a co-insurance arrangement, the assets of the closed block reside

with the reinsurer, addressing some of the credit concerns of the investors. Figure 10.3 shows an example of such a structure. Investors are no longer directly dependent on the credit of the operating insurance company. Even when all closed block assets are transferred to the reinsurer, some of the surplus and additional assets will remain. It is important to note that there are many regulatory requirements that need to be satisfied for this structure to be workable; these requirements depend on the applicable jurisdiction.

Even the use of the more sophisticated structure such as the one presented in Figure 10.3 does not eliminate the risk to investors. Certain risks always remain. Closed block assets might turn out to be inadequate to cover its liabilities. This can happen for a variety of reasons, including initial misestimation of the value of the required assets, poor investment performance or the unexpected increase in the liabilities of the closed block. In addition, policyholder dividends present another element of uncertainty that can also affect the timing of cashflows, possibly jeopardising some of the coupon payments to investors or even repayment of the principal.

It is important to note that investor risk can be significantly reduced if residual risks are reinsured to a non-affiliated reinsurer. The two main risks are longevity and lapsation. These risks have a direct effect on the performance of the closed block and the dividends paid to the policyholders, in turn potentially affecting payments to debtholders.

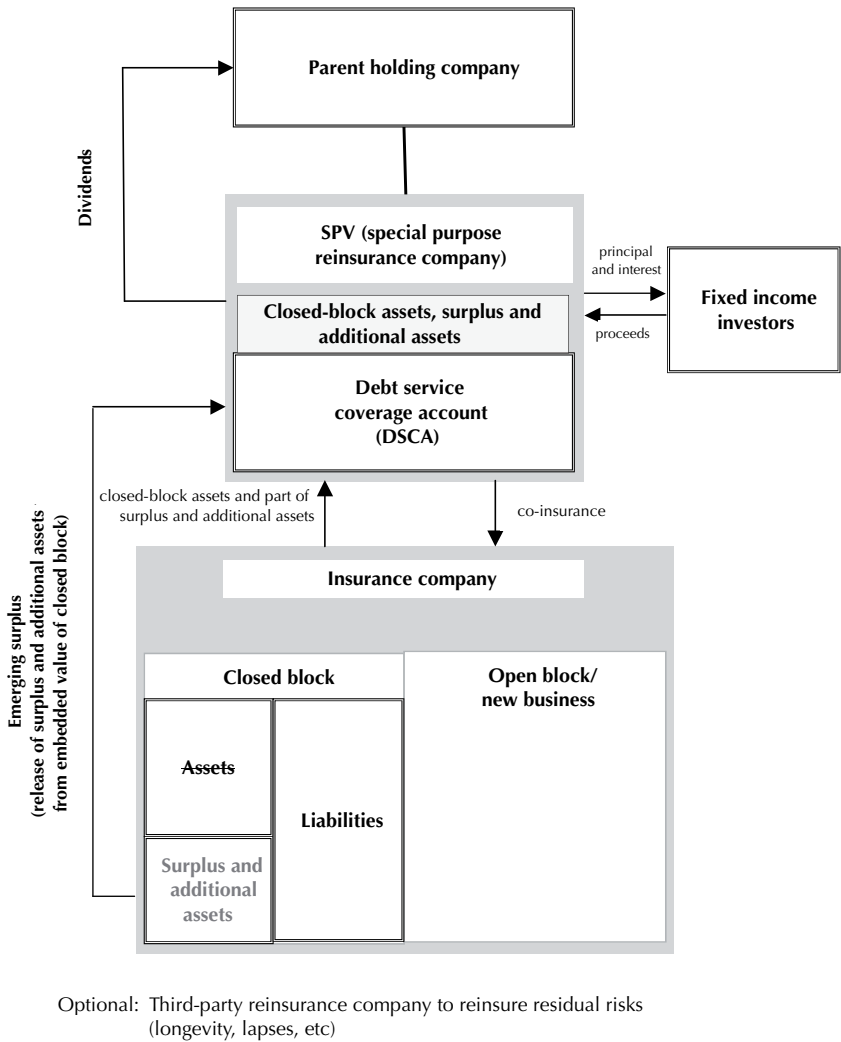
Depending on the type of insurance policies included in the closed block, the tenor of the notes issued to investors could differ significantly. The tenor can be as long as 25 years and in some cases even longer. However, there are artificial ways to reduce the term, such as securitising a smaller part of the closed block. Another alternative is to tranche the securities so that some of them have shorter and others longer tenor.

MODELLING

EV securities are often complex and difficult to analyse, especially for investors unfamiliar with the underlying insurance risk. This difficulty is one of the reasons why credit wraps were so commonly used in the past, when monoline financial guarantors were willing to provide such protection at a relatively low cost.

Actuarial modelling for EV securitisation and monetisation is usually performed by third-party, independent consulting firms. At the very least a third-party firm would provide a comprehensive review of the internally performed analysis. Standard actuarial modelling techniques are traditionally employed. For life insurance, standard modelling software groups

Figure 10.3 Modified embedded value securitisation structure: use of a special purpose reinsurance company



policies by characteristics such as sex and issue age. Each group, referred to as a cell, includes a number of policies for which premium levels, policy benefits, cash value and other policy data are available. The type of data depends on the type of insurance products being modelled. The total number of cells depends on the precision level of the calculation: when lower precision is allowed, some cells are combined. Combining cells and

not using every single cell-specific parameter is the most common approach, driven to a large extent by the limited availability and credibility of data.

The overall approach differs little from the way better insurance companies model their business when no securitisation is involved. Multiple scenarios are generated based on the assumptions in the model, typically with the use of standard software for modelling life insurance. The ultimate goal in this case is to develop a cashflow model for the block of policies and for the investors in the notes; the model should be based on solid assumptions and reflect various scenarios.

Assumptions

Specific assumptions are made for each of the cells being modelled. The first is that of applicable mortality rates and tables. A more technical description of how mortality tables are constructed and used is presented in other chapters. As discussed there, mortality rates are a function of parameters such as current age, age at issue, sex, smoker status, underwriting risk category, type of life insurance policy, face value and others. Each cell has a set of mortality rates based on its characteristics. The mortality tables used in the modelling can span the range between those based entirely on industry experience and those based entirely on the company experience. In most cases, a weighted average of the two is used, with the weight assigned to the company-specific experience being a function of the credibility level of this experience. Some of the larger life insurance companies have accumulated mortality data that has a high level of credibility. A third-party consulting firm has to validate the mortality assumptions by performing a mortality study of the company's actual experience and explicitly taking into account the credibility level of the experience data. Often, the third-party consulting firm will limit this analysis to reviewing the mortality study already performed by the company and will make any necessary adjustments.

Lapse rates represent another important parameter that has to be assigned to each cell based on its characteristics, as is done for mortality rates. Lapses are treated in greater detail in other chapters. It is worth noting that historical lapse experience for a company is not often representative of what the lapse rates will be in the future. Reliable figures for lapse experience for the industry are not available, further complicating the modelling process. As described in other chapters, certain life insurance products can be supported by lapses, in the sense that, without policies lapsing at a certain rate, the overall profitability can be lower than that acceptable to the insurance

company, or even negative. In addition, life settlements can decrease lapse rates while at the same time affecting mortality rates.

Credited rates and minimum guaranteed rates that are part of some life insurance products are another important input. They are based on assumptions, related to the level of interest rates, that may or may not be correct.

Many life insurance products have some optionality embedded in them. Options include taking a loan against the policy, converting to a different type of policy, changing premiums paid while also modifying the death benefit, surrendering the policy for its cash value, and others. The existence of the optionality introduces additional uncertainty to the cashflow projections.

Interdependence of the assumptions

Many of the assumptions are interrelated, further complicating the modelling process. For example, mortality and lapses typically have negative correlation: higher-than-expected mortality rates are usually tied to lower-than-expected lapse rates. Another example is that decrease of interest rates can lead to lower lapse rates.

Direct modelling of correlation among the assumptions is very difficult and rarely performed. The data is insufficient to fully reflect the correlation in the projections; so, instead of improving accuracy of the projections, attempts to incorporate correlation in the modelling process could lead to reduced accuracy and unrealistic scenarios. The quality of the model's output is never better than the quality of the data and assumptions used as input.

Cashflow models

Cashflow models are the foundation of the analysis of closed block and EV securitisation or monetisation. In fact, they are the foundation of the analysis of any securitisation. The base scenario is the one receiving most attention; it is the expected scenario based on the chosen assumptions. However, to analyse the risk and to ensure proper compensation for assuming this risk, investors have to pay particular attention to scenarios that diverge from the expected case – especially the scenarios where the cashflows are insufficient to make payments on the notes when the payments are due. Assessing the extent and probabilities of shortfalls gives investors a picture of the risk involved in investing in the notes.

Assumptions regarding mortality and lapse rates are probabilistic in nature and lead to a multitude of possible outcomes over the lifetime of the closed block or the notes linked to its securitisation. These scenarios have

probabilities associated with them, presenting a more accurate picture of potential investment performance.

The degree and probability of divergence from the expected scenario is dependent not only on the model assumptions, but also on the size of the block of insurance policies being securitised. The bigger the block, the lower the volatility resulting from pure statistical fluctuations. Probabilistic models typically capture this effect relatively well.

Uncertainty related to the choice of assumptions is much more difficult to capture. Qualitative adjustments are often necessary to modify cashflow scenarios so that they will reflect this uncertainty. Any qualitative adjustments are themselves a source of uncertainty and potential error.

Non-actuarial risks

Qualitative adjustments are also necessary to reflect factors whose effect cannot be captured by standard models. For example, it is difficult but necessary to quantify the risk of regulatory action that can have a detrimental effect on the investment performance of the securities. In the case of direct monetisation, it is important to quantify the solvency risk to the operating insurance company due to the poor performance of the ongoing (open) block of policies, which might be a function of such qualitative variables as management quality.

STRESS SCENARIOS

The standard way to analyse risk to investors is by analysing shock events and other stress scenarios. This is different from and complementary to the sensitivity analysis performed as part of the modelling; stress scenarios tend to fall outside the range of those generated in sensitivity analysis.

Mortality shocks and their modelling are described in other chapters, in particular in reference to extreme mortality securitisation. An example of such a shock is a pandemic flu that has the potential to increase mortality to levels significantly above those assumed in the base scenario.

Every assumption can be stressed. For example, shocks to lapse rates can have a significant effect on the cashflows, in some cases almost as significant as mortality shocks. While for mortality a shock is always an increase in mortality rates, for lapse rates in some cases both increases and decreases might need to be considered. Lapse rates in some cases can increase or decrease and then remain elevated or reduced; mortality shocks are more likely to be one-time events, with mortality rates dropping closer to their expected level once an event such as a pandemic flu has passed.

Important stress tests are those that simultaneously stress more than one parameter. These are usually intended to represent specific scenarios for which it is possible to foresee how the various parameters might be affected.

Certain shocks can affect not only the block of policies being securitised or monetised but also the broader insurance market. A mortality shock brought about by a pandemic flu is just one example. Such shocks can affect investors in unexpected ways, starting with losses in the ongoing (open) block happening simultaneously with losses in the closed block, and ranging to scenarios where unexpected regulatory action jeopardises timely payments on the notes. Specific changes to the tax code, or more general economic conditions, can have an effect on policyholder behaviour and simultaneously affect important assumptions such as mortality and lapses as well as the utilisation level of any options embedded in the policies.

Stress scenarios provide important information to investors. The difficulty is usually in determining the chances of such scenarios being realised. Significant judgement is involved in assigning probabilities to stress scenarios.

RATINGS OF EV SECURITISATIONS

Ratings assigned by rating agencies are of great importance in EV securitisation because many investors lack the expertise to analyse these securities independently, and also because rating agencies have access to information not available to the investor community.

Conceptually, in assigning a rating to EV securitisation, rating agencies go through the same main steps as in rating any insurance-linked or other securitisation. Probability of loss is estimated based on the cashflow model presented or the one developed by the rating agency. Loss given default (LGD) is also based on the model. More importantly, rating agencies consider the full range of possible outcomes based on the simulation output of the model.

A rating agency would use the model and its output as presented to form its own conclusions. It will perform sensitivity testing based on the model, either directly or by making specific requests to the firm that performed the original modelling. In addition, the rating agency might choose to build its own model or to engage another consulting firm for this purpose.

Analysis of the structure and legal documents is an important element of the rating process. This analysis can unearth risks to investors not contemplated by the structurers and modellers. Investors highly value this analysis.

Stress testing is an important part of the rating process as well. A rating agency can build its own set of stress tests and analyse their impact on the cashflows and overall risk to investors. Regulatory risk plays a vital role in the analysis and the choice of stress scenarios. As rating agencies are well aware, in EV securitisation this risk can in some cases overshadow those that are explicitly modelled using standard actuarial methods.

Ultimately, the rating is assigned based on the standard default tables that are not specific to rating insurance-linked securities. This makes it possible to perform apples-to-apples comparisons across asset classes. The degree of adjustment based on judgement, however, is probably greater for EV securitisations than for the vast majority of other debt, whether the more traditional type of debt or that related to insurance risk.

Rating caps

In the structure of the type presented in Figure 10.2, the rating of the notes would be capped at the financial strength rating of the operating insurance company, unless a credit enhancement mechanism such as a credit wrap is employed. Even if there is tranching, no tranche would be rated above the rating of the operating insurance company. If the structure reduces or eliminates the dependency on the performance and ratings of the operating insurance company, the notes can be judged on their own merit without the above-mentioned cap. True securitisation, as opposed to direct monetisation, is an example of such a structure.

A rating agency can impose another cap on the ratings of unwrapped tranches, that is, based on the probability of their default within a short period of time that does not allow for gradual downgrades as the credit quality deteriorates. Effectively, the cap is intended to prevent highly rated debt from defaulting due to a single event. Such an artificial cap is not based on quantitative parameters such as the actual probability of default, expected LGD, and for this reason is not considered important by some investors. These investors might assign their own “shadow” rating to the bond, based on default probabilities and not on any cap they consider to be artificial and irrelevant to their analysis. Others, however, fully agree with the approach of assigning such caps, since they are averse to sudden defaults in their investment portfolios. The existence of this cap and the divergent views of investors on this issue are not limited to EV securitisation or monetisation; such factors are more important in the analysis of securities such as catastrophe bonds and extreme mortality bonds. This topic is treated more extensively in the chapters on those securities.

Additional rating caps can be imposed by a rating agency. Such caps, unlike the hard caps as described above, can play more of a guidance role and need not be part of the disclosed formal rating methodology.

Surveillance

After a rating is assigned, the rating agency initiates a surveillance process to assure that any changes to the risk profile of the rated debt are analysed and, if warranted, result in an upgrade or downgrade. General review is performed periodically, but any specific event that can affect the ratings triggers a review. Since rating agencies usually also rate the companies securitising their EV, the analysts should be aware of such events.

EXAMPLES OF EV SECURITISATION

There have been many structures and solutions chosen to monetise EV of insurance business. The structures are still evolving and are expected to continue to evolve.

GRACECHURCH/BARCLAYS EV SECURITISATION

Figure 10.4 shows the structure of an EV securitisation that was not performed in the context of demutualisation and did not require the establishment of a closed block in the traditional sense. Instead, the EV of the whole business of a company put in runoff was securitised.

New Barclays Life was a wholly owned subsidiary of Barclays Bank formed through the merger of Barclays Life Assurance and Woolwich Life Assurance. It was not accepting new policies and was engaged only in managing life insurance and pension business already on the books. Barclays Bank put the company in runoff because it made the decision to distribute insurance products of Legal & General Group PLC instead. Barclays Bank chose to securitise the EV of New Barclays Life primarily in order to obtain regulatory capital relief.

A special purpose reinsurance company, Barclays Reinsurance Dublin Ltd, was formed in Ireland for the sole purpose of providing reinsurance to New Barclays Life. This reinsurance improved the solvency margin of the New Barclays Life and allowed it to repay £752 million in contingent loans to Barclays Bank. The mechanics of the transactions were as follows. Unit-linked assets worth £752 million were transferred from New Barclays Life to Barclays Bank while Barclays Bank transferred these assets to Barclays Reinsurance Dublin Ltd in exchange for an interim bridge loan. Securitising EV of the life insurance business permitted partial refinancing of the loan. It

allowed Barclays to raise £400 million and correspondingly reduce its loan exposure and increase regulatory capital. Another entity, Gracechurch Life Finance PLC, issued £400 million in senior notes with the term of 10 years and also took a subordinated loan in the amount of £352 million from Barclays Bank.

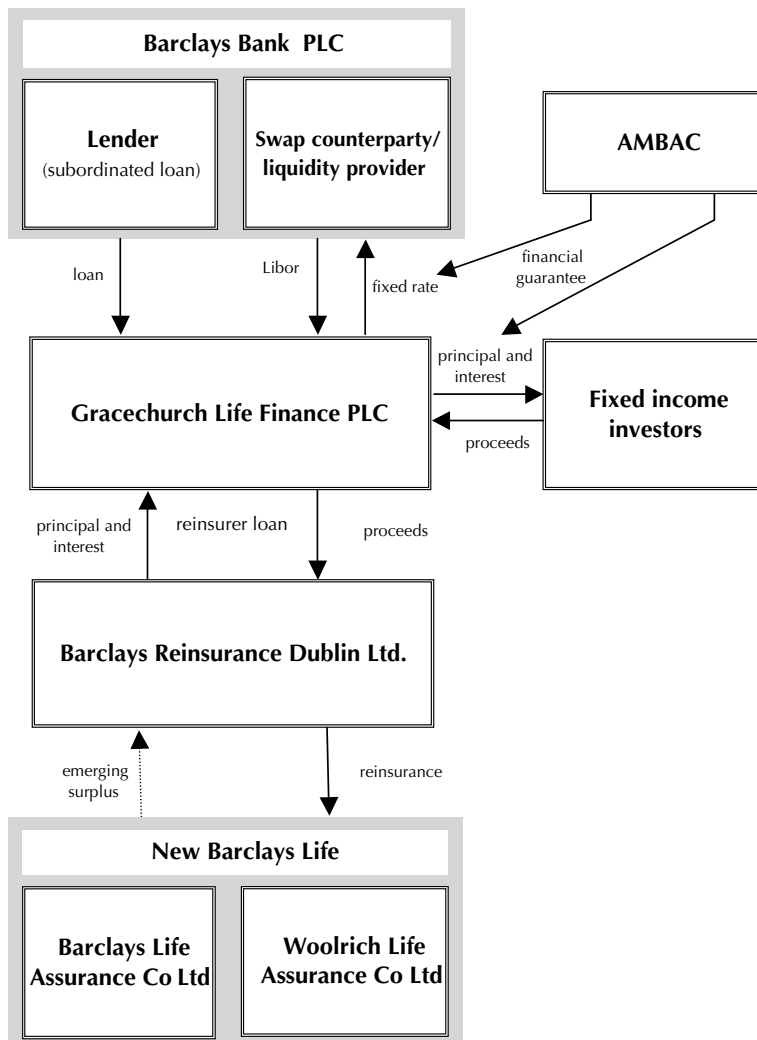
Gracechurch Life Finance PLC made a loan of £752 million to the special purpose reinsurance company, Barclays Reinsurance Dublin Ltd. The transaction was structured so that the SPV reinsurer would repay the loan using the surplus emerging from the business of New Barclays Life. Gracechurch, in turn, would then be able to repay the notes and, once the notes have been repaid, pay back the subordinated loan. Barclays Bank paid the expenses of structuring and executing the transaction.

Barclays Bank PLC played several roles in the transaction. In addition to the ones mentioned above, it served as an interest rate swap provider to exchange the fixed interest paid on the reinsurer loan for the three-month sterling Libor rate, since the senior notes were issued with a coupon tied to Libor. Barclays Bank also served as a liquidity provider. The size of the liquidity facility was set to cover at least two years of interest payments on the senior notes and other payments.

This particular transaction had three types of financial guarantee provided by AMBAC through its UK subsidiary. The main financial guarantee covered timely payments to investors in the notes, including both principal and interest. In addition, AMBAC guaranteed the fixed leg obligations under the interest swap agreement between Gracechurch and Barclays Bank, and the obligations of Gracechurch under the liquidity facility provisions.

The notes were structured to have low risk to investors. The unwrapped rating was A- from Standard & Poor's, while the wrapped rating was dictated by the credit rating of the financial guarantor, which was AAA at the time of issuance. The relatively low risk was a function of the following primary considerations:

- ☐ the credit wrap provided significant credit enhancement;
- ☐ the notes were senior to the sizable subordinated loan, supplying a safety cushion in case surplus would not emerge as projected, resulting in a shortfall;
- ☐ all policies were non-participating and had minimal guarantees;
- ☐ the arrangement was such that it would withstand lower-than-expected investment returns that could reduce the emerging surplus, barring a shock event;

Figure 10.4 Gracechurch Life Finance PLC securitisation structure

- ❑ the mortality assumptions were analysed extensively and judged prudent, including a certain safety margin;
- ❑ the lapse assumptions were also judged to be prudent based on historical persistency data;
- ❑ strong management and support by the parent, Barclays Bank, further reduced the risk; and
- ❑ stress tests were performed by Barclays, rating agencies, and AMBAC.

The Gracechurch transaction remains a reference point in structuring EV securitisations since it provides an efficient way to monetise future surplus emergence.

Scottish Equitable EV monetisation (Zest)

A very different example of EV monetisation is the Zest transaction done by AEGON's Scottish Equitable in 2008. It was a bilateral transaction between Scottish Equitable and Barclays. Barclays then transferred the risk off its balance sheet to investors, but, from the point of view of Scottish Equitable, it was a private bilateral bank transaction.


The structure, presented in Figure 10.5, is a contingent loan made by Barclays against future surplus emergence from a portion of the portfolio of policies held by Scottish Equitable. It had several features that distinguish it from a comparable Portofinos private placement deal completed by AEGON in 2007. Similar to most EV securitisations, it is not a true securitisation but rather a monetisation of VIF.

The interest rate on the contingent loan was fixed and based on the six-year swap rate at the inception. The loan was to be repaid with the surplus emerging from a block of business. The definition of surplus was not standard but rather model-based. The surplus, as defined for this transaction, is insulated from the expense risk, since that risk is retained by the company.

Revolving defined block (RDB) technology was used to define the policies to provide cashflows for the loan repayment. No closed block was set up.

The RDB established against the contingent loan, as implemented in the Zest deal, started with a block of unit-linked policies with the aggregate duration of 15.5 years. Each year for the first three years, Scottish Equitable can change the composition of the RDB by putting additional policies in this block. Additions to the RDB could be policies that, when the defined block was established, were on the books but chosen not to be included in the block. Alternatively, these could be newly written policies. Subject to certain constraints, Scottish Equitable has discretion in what policies to contribute to the revolving defined block during the revolving period and whether to do it at all.

The loan is repaid over the term of the contract. However, over the revolving period emerging surplus can be retained by Scottish Equitable if the amount of surplus in the RDB, with the additional policies added, remains at sufficient levels. In particular, the surplus level in the block should not decline below the base-case surplus, nor should the VIF in the block drop below a certain percentage of the base-case VIF. In the Zest trans-

Figure 10.5 Scottish Equitable VIF monetisation (Zest)

action, this percentage was set at the 95% level. Effectively, over this time the RDB is assured to have sufficient collateral. When the revolving period is over, the loan is repaid through the surplus emerging from the defined block. This deferral of loan repayment is allowed only if the RDB value stays above specific levels.

The notes can be repaid faster if surplus emerges at a greater pace. The stated maturity is 15 years; however, it is expected that the loan will be paid back much faster, most likely in eight years. If the surplus emerging from the defined block is insufficient to meet obligations under the contingent loan agreement, the notes would be written off. This risk always remains.

The risk comprises two primary elements, one having to do with the policies in the defined block and the others with the structure. The main risks for the unit-linked contracts involved are: the investment risk that the returns would be lower than projected; the persistency risk that the lapse rates would be higher than expected; and the risk that the paid-up policy rates would increase beyond expectations. It is important that the three risks be correlated; for example, low investment returns are likely to increase lapse rates. General volatility of the cashflows is also reflected in the modelling. These risks are taken into account in analysing the structure and exploring various cashflow scenarios. In addition, legal agreements and regulatory risk are significant components of the overall analysis. Stress testing played an important role in the evaluation of the investment risks of this transaction.

Zest did not involve any financial guarantee and in this sense is likely to be representative of future EV monetisations for years to come, since the monoline financial guarantee companies are no longer likely to provide this kind of protection, at least not at the cost acceptable to the issuers and investors.

While Zest was a private transaction, it received a rating from Fitch. The

A rating for the Zest VIF notes was subsequently placed by Fitch on negative watch, and later affirmed; but the negative watch had to do more with the financial condition of AEGON and its rating downgrade than with Zest itself. It is not a negative reflection on the structure.

Zest was a private bilateral bank deal with Barclays, but in reality it was then distributed, through Barclays, to investors. Due to the need for the investor to understand the structure and the analysis, it attracted only a small group of investors with the expertise to analyse this transaction and properly assess its risks. It is possible that some investors made their decisions based almost entirely on the Fitch rating and not on their own analysis.

Raising £250 million, the transaction was relatively small for AEGON, but it was significant in introducing some innovative features and, since the £250 million qualified as Tier 1 capital, in providing regulatory capital relief.

TRENDS AND EXPECTATIONS

EV securitisation has been around for a long time. The transaction volume has been growing steadily but rather slowly. There is an expectation that the volume will increase, possibly significantly, as a result of new regulatory developments and the improvement of existing structures. Several factors can contribute to the growth of this type of insurance securitisation.

- ❑ Greater transparency and information availability are key to the continuing growth and development of this market. While the same can be said about all types of insurance securitisations, EV securitisations are some of the least transparent, and many investors in the past have based their decisions entirely on the ratings assigned by rating agencies and on the credit wraps provided by financial guarantors. Investor ability to independently evaluate the securities is a prerequisite for the growth of this market.
- ❑ Transparency should also be extended to the composition of the assets in the special purpose vehicles used for issuing the notes. This will address credit concerns that now permeate the financial industry when any collateral-type structure is involved.
- ❑ Simplifying the securitisation and monetisation structures would make it easier for investors to perform their analysis, while at the same time making it easier to issue the securities and reduce the associated expenses.
- ❑ Reducing credit risk will increase the universe of potential investors and minimise one of the important concerns.
- ❑ Shortening maturities of the notes overall and tranching the debt so that

it includes shorter maturities will also serve to increase the universe of potential investors and contribute to the growth of this market.

- ❑ Regulatory changes, making it easier to securitise EV and removing some of the regulatory risk described above, can lead to the growth in such transactions.
- ❑ Rating agencies' becoming more comfortable with EV securitisations can also make investors more comfortable with these securities and make it easier to execute such transactions.
- ❑ Improved modelling would give investors greater confidence. Better disclosure of modelling results, including assumptions and sensitivity analysis, would increase the level of confidence even further.
- ❑ Broad regulatory developments currently in motion, in particular in Europe, can have a sweeping effect on the way insurance companies manage their capital and risk. Solvency II is one such important development. Securitising EV is a way of managing capital and risk, and these regulatory developments are expected to lead to new transactions of this nature.

We are witnessing the growth of securitisation or monetisation of EV not limited to the context of demutualisation, leading to EV securitisation becoming part of the capital management toolkit for insurance and reinsurance companies. Solvency II can become a catalyst of this process.

Investors will grow in their sophistication and the ability to analyse these securities. One of the by-products of this process will eventually be lower returns demanded by investors for the same level of risk, which will be more in line with other securities. This, in turn, will make EV monetisation more efficient for issuers.

Finally, it is expected that securitisation of future cashflows from other types of insurance risk, not necessarily life insurance, will grow. So far, few such securitisations have been executed.

¹ Technically, the use of this mechanism does not always meet the standard definition of over-collateralisation. However, it is sufficiently similar to use this term in the context of securitising future cashflows from insurance business.

Part IV

Investing in and Modelling Securities Linked to Mortality and Longevity Risk

Securitisation of Extreme Mortality Risk

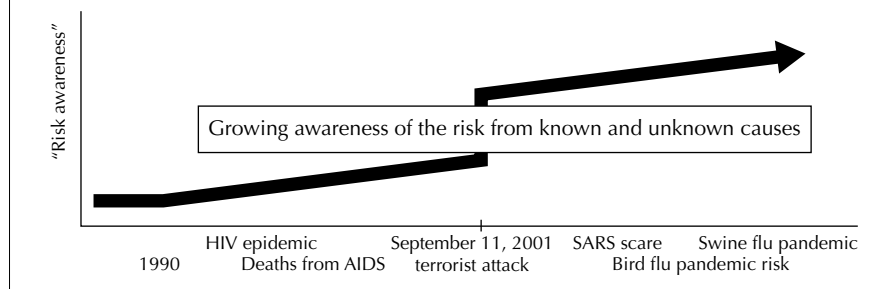
This chapter examines extreme-mortality risk and how this risk is transferred to the capital markets through securitisation. It describes extreme-mortality bonds and their basic structures. While modelling mortality risk is explained later in this section, this chapter provides guidance on the current modelling approaches for extreme mortality, which can be of value both to (re)insurance companies who want to transfer this risk to the capital markets and to investors in extreme-mortality bonds.

THE RISK OF EXTREME MORTALITY

Mortality is integral to life insurance and annuities, with mortality rates being a key component of setting both price and reserve levels for a life insurance company. To a significant degree, life actuarial science is focused on analysing mortality and producing mortality tables. Mortality rate measures the number of deaths in a period of time in a population relative to the size of that population. The “population” could differ from the general population; it can be, for example, age- and gender-specific. Insured populations usually exhibit mortality characteristics different from those of the general population. Mortality rates tend to be stable or exhibit easily identifiable trends. Insurance companies, having large portfolios of life insurance policies, take comfort in this stability.

Mortality risk is the risk that actual mortality will turn out to be greater than projected. This risk is assumed by companies writing life insurance. The reverse of mortality risk is the risk of longevity, that is, of people living longer than expected and longer than assumed in the estimation of financial liabilities. Insurance companies writing annuity products are subject to this risk, as are pension funds providing defined benefits to participants.

Historically, mortality risk was not considered particularly important by insurance companies because of the relatively high predictability of mortality rates for large pools of insured lives, as well as steady declines in

Figure 11.1 Timeline of the growing awareness of the risk of extreme mortality

mortality due to people living longer. This view has been slowly changing over the years. The change started with the HIV/AIDS epidemic, which highlighted the risk of sudden increases in mortality rates. The events of September 11, 2001, brought additional attention to the risk of mortality shocks. A jump in mortality rates could be caused by a terrorist attack or an event such as a flu pandemic. It is likely that the risk of such sudden jumps has been increasing; even more importantly, awareness of the existence of this risk has been growing.

Realisation that sudden increases in mortality rates represent a significant risk to insurance and reinsurance companies has led to the growing demand for reinsurance protection against this risk. Concurrently, life reinsurance companies, who act as aggregators of risk, have become more aware of the risk concentration in their portfolios and less willing to provide this type of protection. Especially after the events of September 11, 2001, reinsurance of extreme-mortality risk has become very expensive. Traditional reinsurance now often excludes catastrophic events. The situation parallels the “Katrina effect” in the property insurance industry, albeit on a smaller scale. The H1N1 2009 pandemic and the general swine flu scare continued to bring attention to the risk of significant spikes in mortality rates.

SECURITISATION OF EXTREME MORTALITY RISK

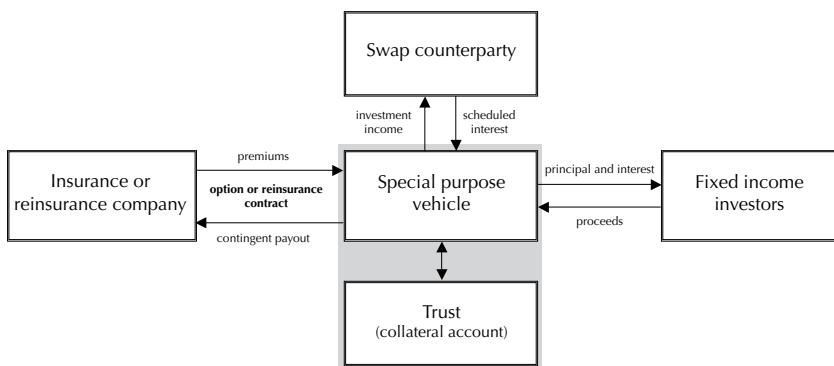
In some cases the risk of extreme mortality is lessened because the same insurance company is writing both life insurance and annuity products, and the increase in life insurance claims can be partly offset by the decrease in annuity liabilities. Overall, however, the problem of risk accumulation in the life insurance industry is real and in need of a resolution.

Transferring some of the risk of extreme mortality to the capital markets

is a natural solution to the problem. Capital markets are much larger than the life insurance industry and in a better position to absorb this risk. As long as the potential risks and returns of a financial instrument can be quantified, capital markets participants will be willing to invest in it. Such an instrument for transferring the risk of extreme mortality could be structured in the form of a fixed income security similar to a property catastrophe bond. (Property catastrophe bonds are described in Chapter 3, which also provides a more detailed treatment of the structuring mechanics for such securities.) Figure 11.2 shows a generic structure of an extreme mortality bond, with the insurance or reinsurance company sponsoring the bond entering into a contract with a special purpose vehicle (SPV) to be reimbursed for losses due to extreme mortality events. Simultaneously, the SPV issues fixed income securities to investors, with the repayment of principal and payment of interest tied to the occurrence of the same extreme mortality events specified in the agreement with the insurance company. This agreement can be in the form of an option or reinsurance contract. In the latter case the SPV is a special purpose reinsurance company.

The structuring mechanics are similar to those used for property catastrophe bonds, with the exception that a mortality-based index is created to act as the bond default trigger. The similarity extends to the swap counterparty, collateral account and other credit-risk issues that came to light in the aftermath of Lehman Brothers' bankruptcy and necessitated changes to the standard structures. Below we examine the structure used in the first extreme-mortality bond ever issued, Vita Capital.

Figure 11.2 Typical extreme mortality bond structure



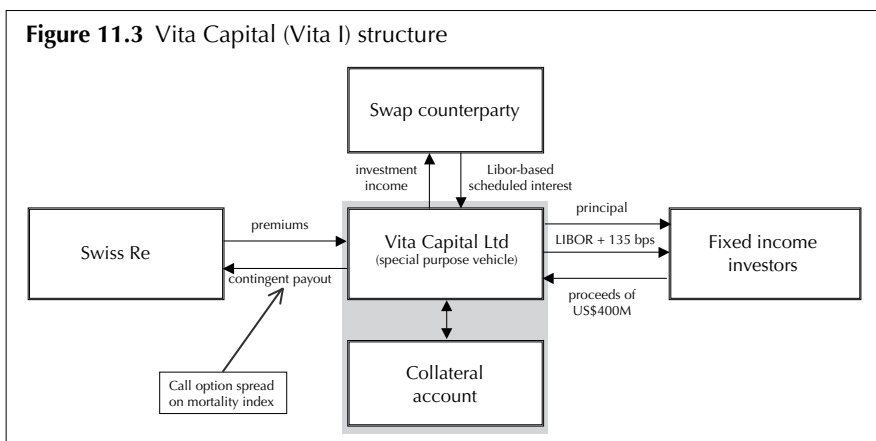
THE GROUNDBREAKING VITA SECURITISATION

Transferring extreme mortality risk directly to the capital markets is a solution first used in 2003 by Swiss Re in the Vita transaction. Swiss Re's approach echoed that of a catastrophe bond issuance repeatedly employed for the risk of natural disasters.

The Vita Capital transaction was the first securitisation of extreme mortality risk. Sponsored by Swiss Re in December 2003 and maturing in January 2007, the catastrophic mortality bond was structured to reduce the exposure of Swiss Re to a sharp increase in mortality. The total issue size was US\$400 million. (The size of the transaction as initially presented to investors was US\$250 million; a follow-up US\$150 million issue was planned for the next year. The unexpectedly strong investor demand allowed Swiss Re to combine the two issues.) The trigger was a weighted average of the general-population mortality rates in five countries: the US, the UK, Italy, France and Switzerland. The index was constructed to reflect the exposure of Swiss Re's life insurance book to adverse mortality experience in these five countries. It is likely that the choice of the countries was influenced, at least in part, by the availability of reliable government data on population mortality.

Structure of the Vita Capital transaction

Vita Capital Ltd, an SPV, was established for the securitisation. The SPV simultaneously entered into the following two transactions. The first was an agreement with Swiss Re to provide, in exchange for a premium paid by Swiss Re, a call option on the SPV assets. The option trigger was tied to a



population mortality index. The second transaction was issuing a fixed income security to investors.

As in one of the typical property cat bond structures, returns from the collateral account were swapped for a Libor-based rate with a highly rated counterparty. This reduced the interest rate risk and made the bonds floating rate instruments. The assets of the SPV were invested in high-quality financial instruments.

The option contract between Swiss Re and Vita Capital was in the form of a call option spread on a mortality index. The lower strike price, that is, the start of payments to Swiss Re from Vita Capital, was set at 130% of a specified value of the index. The upper strike price, leading to full payment to Swiss Re, was set at 150% of the same value of the index.

Trigger index

Designing the right type of index for this pioneer transaction was a difficult task. The first objective in designing the index was the minimisation of Swiss Re's basis risk. As much as possible, the index was supposed to mimic the actual exposure of the company to the extreme mortality risk. The second objective was to use verifiable data and achieve the greatest degree of transparency for investors.

Figure 11.4 shows the distribution by country in the index. Only government sources were used for obtaining mortality data. All five are developed countries with relatively stable mortality patterns for general population. The government data-reporting agencies (or their predecessors) in these countries have a long track record and expertise in data collection.


The distribution by gender used in the index roughly corresponds to the likely gender distribution in the actual life reinsurance portfolio of Swiss Re: 35% female and 65% male.

Table 11.1 shows the distribution by age within the index. This age distribution is not atypical for a diversified life insurance portfolio and is likely close to the actual distribution of the portfolio of Swiss Re.

The index value was calculated as a weighted-average mortality rate, with averaging over country, gender and age based on the weights specified

Table 11.1 Distribution by age within the Vita index

Age group	20–24	25–29	30–34	35–39	40–44	45–49	50–54	55–59	60–64	65–69	70–74	75–79
Weight	1%	5%	12.5%	20%	20%	16%	12%	7%	3%	2%	1%	0.5%

Figure 11.4 Geographic distribution within the Vita index

Note: The following sources were used for population mortality data reporting:

- US: Centers for Disease Control and Prevention, National Center for Health Statistics
- UK: Office for National Statistics
- Italy: Istituto Nazionale di Statistica
- France: Institut National de la Statistique et des Études Économiques
- Switzerland: Swiss Federal Statistical Office

above. The base value of the index to serve as a comparison point was chosen to be the 2002 mortality level.

Payout schedule

The payments to Swiss Re, corresponding to the reduction of principal repayment to investors in the bond, occur when the index value exceeds 130% of the base value and increase proportionally until it reaches 150%, at which point the full amount is owed to Swiss Re and investors receive no principal repayment. Figure 11.5 shows the reduction of principal repayment to investors based on the value of the index.

In this first Vita deal, a one-year calculation period for the index was used.

Benefits of the Vita transaction to Swiss Re

The transaction allowed Swiss Re, the world's largest life and health reinsurance company, to protect itself against the risk of a catastrophic mortality event. It contributed to the more efficient use of capital by the company by reducing the economic capital required to support its book of reinsurance business. It had a positive effect on the company's regulatory capital

PANEL 11.1 REFERENCE INDEX CONSTRUCTION

Following (with some modifications) the notation used by Cummins (2004), we could write a general formula for index construction as

$$q_t = \sum_j C_j \sum_i A_i \left(G^{male} q_{ijt}^{male} + G^{female} q_{ijt}^{female} \right)$$

where q_t is the value of the mortality rate index based on the data reported as of time t (or based on data from period t),

q_{ijt}^{male} is the mortality rate for male lives in age group i in country j ,

q_{ijt}^{female} is the mortality rate for female lives in age group i in country j ,

C_j is the weight assigned to country j ,

A_i is the weight assigned to age i ,

G^{male} is the weight assigned to males, and

G^{female} is the weight assigned to females

The value of the index could be compared to the base value of q_0 . In the case of Vita Capital, the base value of the index is that of the year 2002, that is, q_{2002} .


An even more general formula for index construction could include age weights differing by country and by gender, as well as male/female distribution varying by country. In this case, the formula for index construction becomes

$$q_t = \sum_j C_j \sum_i \left(A_{ij}^{male} G_j^{male} q_{ijt}^{male} + A_{ij}^{female} G_j^{female} q_{ijt}^{female} \right)$$

If the value of the q_i index is intended to represent an actual mortality rate, care should be taken to ensure that it is scaled appropriately.

requirements and on the capital requirement imposed by rating agencies in order to maintain high credit ratings.

Swiss Re also benefited from the collateralised nature of the mortality bond. Unlike the use of reinsurance to transfer risk, the Vita transaction did not expose Swiss Re to the credit risk associated with the creditworthiness of the reinsurance transaction counterparties. This risk would be particularly high in the event of significant overall increases in mortality rates.

Figure 11.5 Reduction in principal repayment to investors based on the Vita index

OTHER SECURITISATIONS OF EXTREME MORTALITY RISK

The Vita Capital transaction, which is now referred to as Vita I, was the harbinger of a number of other extreme mortality risk securitisations. The transfer of extreme mortality risk to the capital markets is expected to continue to grow even though the growth to date has been uneven. Table 11.2 shows some of the extreme mortality bonds issued.

The basic structure of the bonds has not changed, even though some new elements have been added to make the bonds attractive to a wider universe of investors and to make the structure more efficient for the issuer in terms of basis risk and other considerations.

Table 11.2 Extreme mortality securitisations

Company	Year	Principal amount	Number of tranches
Swiss Re – Vita Capital	2003	US\$400 million	1
Swiss Re – Vita Capital II	2005	US\$362 million	3
Scottish Re – Tartan Capital	2006	US\$155 million	2
AXA – Osiris Capital	2006	€150 million and US\$250 million	4
Swiss Re – Vita Capital III	2007	€240 million and US\$390 million	2
Munich Re – Nathan Ltd	2008	US\$100 million	1
Swiss Re – Vita IV	2009	US\$75 million	1

Note: The AXA deal included B1 and B2 tranches that were identical in all terms, with the exception of B1 having been wrapped by an AAA-rated financial guarantee company and consequently having a higher credit rating.

One key development is the slicing of risk into tranches with different risk–reward characteristics. For example, the Tartan Capital transaction includes two tranches with different attachment and exhaustion points. The riskier tranche, Class B notes, had the attachment point of 110% of the index used in this transaction, with the exhaustion point of 115% of the index. If Class B notes were to suffer full default, the less risky tranche would be activated with attachment point of 115% and exhaustion point of 120% of the index.

Another development was the use of credit wrap to enhance the ratings of a specific tranche and make it attractive to a broader group of investors. In the Tartan Capital transaction, Class A notes were wrapped by an AAA-rated financial guarantor. This resulted in Class A notes being rated AAA/Aaa compared with the BBB/Baa3 rating for Class B. The changes in the financial landscape have led to such financial guarantees being no longer available, at least at a reasonable cost, and it is unlikely they will be incorporated in any future extreme mortality bond structures.

The index in the Tartan securitisation was chosen with different age weights for males and females, to more accurately replicate the actual insurance portfolio of the bond sponsor, Scottish Re. The calculation period was

PANEL 11.2 MORTALITY RATE DEFINITION FOR THE TARTAN TRANSACTION

Mortality rate (not scaled because the sum of the weights is not 100%) for year t is defined as

$$\tilde{q}_t = \sum_i (A_i^{\text{male}} G^{\text{male}} q_{it}^{\text{male}} + A_i^{\text{female}} G^{\text{female}} q_{it}^{\text{female}})$$

The scaled value, representing the true mortality rate, is


$$q_t = \sum_i (W_i^{\text{male}} q_{it}^{\text{male}} + W_i^{\text{female}} q_{it}^{\text{female}}),$$

where W_i^{male} and W_i^{female} are the actual weights applied to age group i for males and females respectively.

The Mortality Index Value, calculated over a period of two consecutive years, is then expressed as

$$Index_t = \frac{\frac{q_{t-1} + q_t}{2}}{\frac{q_{2004} + q_{2005}}{2}} \times 100\%$$

The value of the index for the 2004–2005 24-month period is used as the base.



selected to be two years, as opposed to the one year in the original Vita Capital transaction.

The index measurement period for the transaction is illustrated in Figure 11.6.

As in the Vita deal, in the Tartan transaction a measurement period might show a loss, with the loss increasing linearly between the attachment and exhaustion points of the index. If both measurement periods show losses, the greater of the two loss percentages is chosen to determine the ultimate loss amount used in the transaction settlement. Panel 11.3 shows how the loss percentage is calculated.

PANEL 11.3 DEFINITION OF LOSS TRANSACTION FOR TARTAN CAPITAL

The loss percentage for this transaction was defined as

$$\text{Loss percentage}_i = \frac{\text{Index}_i - \text{Attachment}}{\text{Exhaustion} - \text{Attachment}} \times 100\% \quad \begin{array}{ll} 0\% & \text{for } \text{Index}_i \leq \text{Attachment} \\ 100\% & \text{for } \text{Attachment} < \text{Index}_i \leq \text{Exhaustion} \\ & \text{for } \text{Index}_i > \text{Exhaustion} \end{array}$$

The use of a calculation period longer than one year is expected to become standard in future extreme mortality securitisations. It also appears that large insurance and reinsurance companies are willing to use an index based on general population mortality data, even though this introduces basis risk.

BASIS RISK

In extreme mortality risk transfer, the issue of basis risk is greater in importance than in many other types of insurance securitisations. In the mortality bonds issued so far, the trigger has been based on the government population indexes instead of on the actual mortality of insured policyholders. This approach is favoured by investors for its transparency and the elimination

of the need to examine underwriting standards of the insurance company, which is something that often cannot be done well even by experts in life insurance underwriting. While this type of default trigger is favoured by investors, it creates the problem of basis risk for the insurance or reinsurance company transferring the risk to the capital markets.

Basis risk measures the chances that the bond will prove to be an ineffective hedge, with the company suffering extreme mortality losses while the bond is not triggered. Another undesirable scenario is that the bond is triggered while the issuer or sponsor has not experienced catastrophic mortality losses. Careful structuring of a mortality bond minimises the likelihood of both scenarios.

Matching mortality experience of a block of insurance policies to a general population index for extreme mortality securitisation is different from the basis risk analysis performed in the Regulation XXX or embedded value securitisation context. Currently, the main risk in catastrophic mortality bonds is believed to come from a pandemic of swine flu or a similar disease. In such a pandemic, the segments of the general population most likely to be severely affected are children and the elderly. Mortality experience of these two segments is likely to be the driver of the general-population mortality index in this scenario. However, these two segments are usually less likely to be in the pool of insurance policies than in the general population. The end result of this mismatch is that in a pandemic the mortality experience of the insured lives is likely to be significantly better than that of the general population.

While swine and bird flu are currently considered to be the main potential sources of extreme mortality, other diseases could affect a different segment of the population. The HIV/AIDS epidemic, while not resulting in the huge loss of life that was initially feared, is an example of a risk affecting the segment of the population likely to be sufficiently represented in the insured pool. Unlike the flu, the HIV virus has affected primarily adults not in the elderly category. Another deadly disease, should one emerge, could also disproportionately affect a specific segment of the population.

CREDIT ENHANCEMENT

In the past, in extreme mortality securitisations, credit enhancement was often accomplished by adding a credit wrap to the securities. Such a credit wrap was generally provided by a monoline financial guarantee company. The credit wrap added value to the transaction by significantly expanding the investor base as well as enhancing liquidity. It also provided a certain

degree of comfort to investors, since financial guarantors generally either performed an independent analysis of the risk or validated the analysis already performed, reviewed the integrity of the overall structure and carefully examined the documentation. Even investors buying unwrapped tranches received some assurance from the fact that a financial guarantor had analysed the structure and documentation. (Of course, they should perform their own analysis; the unwrapped tranches should be modelled differently from the wrapped ones, and numerous additional considerations are involved.) Financial guarantors also assumed some of the residual risks embedded in the securities. Due to the reduced risk, wrapped bonds commanded a lower spread.

The credit crisis made such financial guarantee unavailable, but even before that, concerns had been raised that financial guarantors would have limited capacity to take on the risk of a single event such as a pandemic. Due to potential “risk stacking”, they might have been unable to provide credit wrap if more and more catastrophic mortality bonds were issued. At the very least, even before the credit crisis we might have expected that credit wrap for extreme-mortality risk could become more expensive.

The changes in the financial markets’ landscape have resulted in financial guarantee no longer being obtainable. It is not expected that the situation will change in the near future; most likely, financial guarantors will never provide this kind of protection at a reasonable cost, and financial guarantee will never be part of the extreme mortality bond structures.

INVESTOR TYPES

The universe of potential investors in extreme-mortality bonds was very large for wrapped tranches. The high ratings afforded through the use of financial guarantee open this class of fixed income instruments to investors who would invest only in very low-risk securities. On the other hand, unwrapped tranches are attractive to a much more limited number of investors. Many investors shy away from these securities because of the novel nature of the risk as well as the difficulty of properly quantifying it. As more of the extreme mortality bonds are issued, investors will become increasingly familiar with these securities and will be more willing to purchase them in the now standard unwrapped form.

EXTREME MORTALITY RISK QUANTIFICATION AND PRICING

In transferring extreme mortality risk to the capital markets, both (re)insurance companies and investors have to be able to model the risk. An investor

has to be able to assess the risk of default of an extreme mortality bond to decide whether to buy this security.

Traditional actuarial approaches to modelling mortality risk are not usable in the context of analysing extreme mortality. These approaches have been developed to use historical data for quantifying stable mortality rates and, to some degree, identifying and incorporating in the rates the trends of slowly shrinking mortality. They remain essential in the context of setting prices and establishing reserves for life insurance policies. However, extreme mortality events, by their very nature, are generally not represented in the data collected by insurance companies. These events are something that has not happened in recent history, nor, indeed, has ever happened in the history of the life insurance industry. A standard mortality table is of little use in trying to quantify the risk of a sudden jump in mortality rates due to an event such as an influenza pandemic or a large-scale terrorist attack.

Modelling mortality rates

The regulatory – and often internal – pressure has long been on setting up mortality tables in a “prudent” fashion that would avoid underestimating the rates. More recently, especially with the growing attention to the efficient use of economic capital, the focus has been shifting from prudent to accurate mortality rates. However, the approach has remained largely deterministic and it is extremely rare to see stochastic modelling of mortality rates in a traditional life insurance setting.

While the life insurance industry has recognised the need to model interest rates in a probabilistic manner, this approach has not yet found its way to the modelling of mortality rates.

Factors affecting mortality risk

The factors that affect mortality risk of a life insurance or reinsurance company could be split into the following four categories (of which the last one is of particular interest in this context):

1. **RANDOM STATISTICAL FLUCTUATION:** Statistical fluctuations are expected and are a function primarily of the size of the block of insurance policies, with larger pools of insured lives showing smaller fluctuations relative to the mean. Reinsurance companies and large primary insurance companies tend to have very sizable pools of lives and to be less affected by random statistical fluctuations than smaller companies. Another factor affecting the impact of random statistical fluctuations is the homogeneity of the policies within the pool.

2. **MISESTIMATION OF GENERAL MORTALITY TRENDS:** In the US and most other countries, mortality rates have been experiencing steady reduction over many years. On average, people live longer, and the historical mortality data is not always directly applicable to a pool of current insurance policies. The projection of the trend into the future, however, is a very difficult task. Without fully understanding all causes of mortality and their change over time, we cannot simply assume that the current trend can be extrapolated into the foreseeable future.
3. **DATA ISSUES AND MISCALCULATION OF CLAIM LEVELS:** Calculating mean expected mortality values could introduce a systematic mistake due to data issues. It is possible that underwriting classes and underwriting standards within each class have been changing or “drifting” over the years, affecting the reliability of historical mortality data used for calculation of mean values. For smaller books of business, random fluctuations in mortality could lead to the misestimation of mortality levels. The effect of longevity improvements over time, if not taken into account appropriately, could also contribute to the miscalculation of mean values.
4. **CATASTROPHIC EVENTS:** Catastrophic events are, by their very nature, difficult or impossible to model based on the traditional data used for estimating mortality rates. They are unlikely to be in the historical data of an insurance company – major events such as the Spanish Flu pandemic of 1918 happened too long ago to be usefully included. The impact of this same event today would likely be quite different from what it was in 1918. In addition, many of the potential causes of catastrophe mortality events are new and by definition cannot be found in historical data. Finally, there are bound to be events that we are not in a position to foresee and model today.

It is important to point out that most probabilistic models of mortality are not well suited to describing extreme mortality events. Approaches such as the Lee–Carter model or those borrowed from interest-rate modelling are very useful in most applications, but they do not easily allow us to model mortality jumps corresponding to extreme mortality events. Models that explicitly include the jump component are very difficult to parameterise based on available data.

CURRENT MODELLING APPROACHES

Modelling of extreme mortality is still in its infancy, with approaches being developed and refined. In the context of securitising extreme-mortality risk, the ability to quantify this risk in a probabilistic framework is of paramount importance. While this ability is important to insurance companies wishing to transfer the risk to capital markets, it is of even greater importance to investors in securities based on the risk of extreme mortality. Investors have to have the ability to assess the risk – both its overall level and its potential correlation with other assets – and determine the level of compensation appropriate for taking the risk, that is, the price of the securities. Uncertainty in the reliability of modelling results leads investors to demand greater return for investing in the securities.

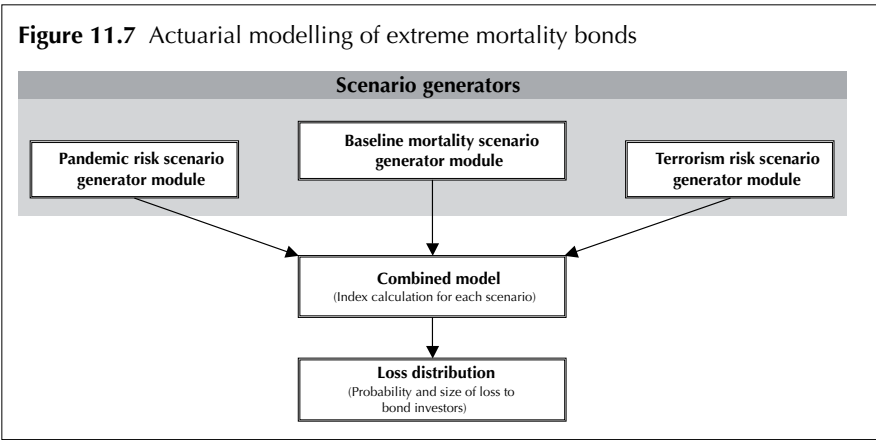
Natural disasters, while capable of causing huge economic losses, have not had a significant impact on mortality rates in the US and most developed countries. Developing countries, in particular in Asia, are more exposed to this risk, but life insurance is less common in these countries and amounts insured are relatively low. For example, while the death toll of the Asian tsunami of 2004 was over 150,000, the vast majority of these people were not insured. The mortality risk to life insurance companies is concentrated in the developed countries and in particular in the US.

The current way of modelling mortality rates in the context of extreme mortality securitisation involves independent modelling of the following three major components of mortality rates:

- ❑ baseline mortality, reflecting statistical fluctuations around the standard mortality mean;
- ❑ terrorism component, which reflects the effect of potential terrorist attacks on mortality rates; and
- ❑ pandemic component, which reflects the effect of large-scale epidemics of serious infectious diseases on mortality rates.

The above components are the only important contributors to extreme-mortality risk. Although war is another obvious driver of population mortality, it is excluded as a cause of death from most life insurance policies and thus has limited effect on the mortality experience of a typical life insurance company. Figure 11.7 illustrates the modelling approach that has been used to analyse extreme mortality bonds.

Independent scenario generators are created for stochastic modelling of each of the major components of mortality rates affecting an extreme mortality index. A large number of scenarios are produced for each of the




three components. These outputs are then combined at the next level of the model, where the total mortality rates and index values are calculated for each scenario combination. This produces a probability distribution of index values, which in turn could be used to determine the probability distribution of losses to investors in the extreme mortality bond.

Below, we take a look at how each of the mortality rate components could be modelled in the framework described above.

Component 1: Baseline

Modelling the baseline component of the mortality rates involves generating mortality scenarios including statistical fluctuations around the expected value. Historical mortality data is utilised; one such approach (used by Milliman, Inc, in providing actuarial analysis to support extreme mortality risk securitisation) is based on time series iterations, similar to the bootstrap



method. A stochastic error term is introduced to model the volatility of mortality rates. Time series are fitted by minimising the sum of squared errors of mortality rates, or by some other method.


Base aggregate mortality exhibits fluctuations that are very small compared with the attachment points for extreme mortality securitisations, with index values based purely on baseline mortality clustered tightly around the mean in the vast majority of scenarios.

Component 2: Terrorism

At present there is no established model of terrorism risk. Attempts to indirectly assess the risk of terrorism, such as through the proposed introduction of financial “terrorism futures”, have not been successful. For lack of a better way, the probability of a terrorist attack is now being assessed based on expert analysis and approaches such as the Delphi method. Assessing mortality resulting from a terrorist attack is even more difficult.

One approach to modelling the terrorism component of the mortality rates is the use of a multilevel logic-tree approach. As utilised by Milliman, Inc, in modelling the Tartan Capital securitisation (Scottish Re), quarterly frequency of terrorist events was based on a normal distribution, with the mean and standard deviation taken from the actual data for 1999–2004 of all terrorist attacks on American citizens and property, excluding events in Afghanistan and Iraq (see Figure 11.9).

Figure 11.9 Frequency and severity of terrorist attacks against the US



Sources: Milliman, Inc; US State Department; National Counterterrorism Center of the Office of the Director of National Intelligence

For determining severity, that is, number of deaths, at each level of the logic tree there were three choices:

- ❑ “success” of the terrorist attack, resulting in a random number of deaths in a predetermined range;
- ❑ “failure” of the terrorist attack (no deaths); and
- ❑ escalation to the next level of severity (greater number of deaths).

Probabilities of each outcome – “success”, “failure” and escalation – at every level were determined by fitting an exponential distribution to the data in Figure 11.9.

This modelling approach is imperfect in its reliance on such limited data, and will certainly be improved in the future. However, while imprecise, it has served the purpose of demonstrating that the terrorism component is not the driving force behind potential significant increases in mortality rates. In fact, in a stochastic framework, the effect of terrorism on mortality rates is small due to the relatively low probability of a large number of deaths from a terrorist attack. While an individual life insurance company might have a concentration of risk in a terrorism-prone location, the effect on the general population mortality index is exceedingly unlikely to lead to truly catastrophic deviations from the mean. (Nuclear attack by terrorists is a possible reason for a catastrophic jump in mortality rates due to terrorism.)

Component 3: Pandemic

Pandemics are the key driver of potential jumps in mortality rates. Outbreaks of serious infectious diseases have the potential to cause many

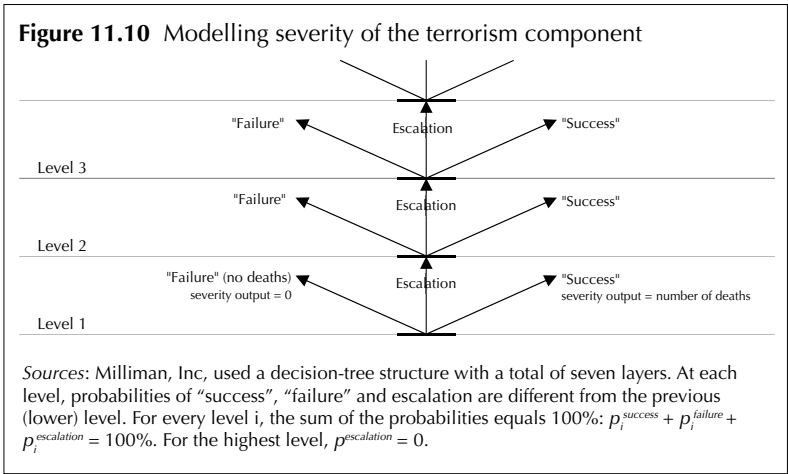


Table 11.3 Flu pandemics and critical development in the last 100 years

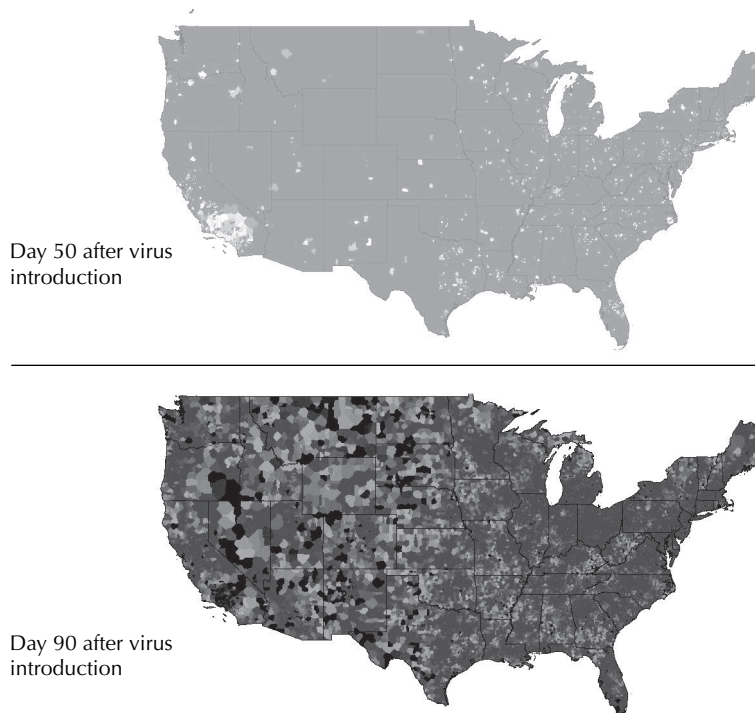
1918	“Spanish flu” H1N1 Over 500,000 deaths in the US (20 million to 50 million worldwide)	Pandemic
1957–58	“Asian flu” H2N2 70,000 deaths in the US	Pandemic
1968–69	“Hong Kong flu” H3N2 34,000 deaths in the US	Pandemic
1977	“Russian flu” H1N1 Appearance of new influenza strain in humans	
1997	H5N1 – first flu virus found to transmit from birds to people Appearance of new influenza strain in humans	
1999	H9N2 – probable transmission from birds to people Appearance of new influenza strain in humans	
2002	H7N2 – possible transmission from birds to people Appearance of new influenza strain in humans	
2003	H5N1, H7N7, H7N2, H9N2 Appearance of new influenza strain in humans; spread of H5N1	
2004	H5N1, H7N3, H19N7 Appearance of new influenza strain in humans; spread of H5N1	
2005	H5N1 Spread of H5N1	
2006	H5N1 Spread of H5N1	
2007	H5N1, H7N7 Spread of H5N1, appearance of H7N7 strain in humans	
2008	H5N1 Spread of H5N1	
2009	H5N1, H1N1 Spread of H5N1, appearance and rapid spread of H1N1 around the world	Pandemic

Source: National Institutes of Health, US Department of Health and Human Services

deaths. Pandemics of bubonic plague in the Middle Ages wiped out a significant portion of the population in many countries. New diseases or new strains of known diseases continue to emerge. Table 11.3 shows flu pandemics and critical developments in the emergence of new strains of the flu virus, including the H5N1 flu strain (bird flu) that has been found to transmit from birds to people. If the virus mutates further and easy human-to-human transmission becomes possible, the result could be a devastating pandemic with a very high death toll. The table also shows the swine flu pandemic in 2009.

Figure 11.11 shows an illustrative scenario of the spread of pandemic flu in the US. It is one of the stochastic scenarios generated by a large-scale simulation model on a supercomputer in Los Alamos National Laboratory. The model examines the rapid spread of a pandemic influenza virus strain in the continental US, starting with the arrival of 10 infected individuals in

Figure 11.11 Illustrative scenario of a pandemic flu outbreak in the US



Sources: Los Alamos National Laboratory, US Department of Energy


Los Angeles. The model attempts to reflect the response of the population to the outbreak, including decreased travel and other mitigation strategies. Averaged over all scenarios, the pandemic peaks about 90 days after the introduction of the virus in the absence of vaccination or antiviral medications.

Historically, pandemics have come in waves; the Los Alamos National Lab supercomputer simulation (Figure 11.11), while very sophisticated, shows only the first wave. Unfortunately, a stochastic model of such complexity cannot be directly used for mortality modelling at this time.

An approach currently used in modelling the pandemic component of mortality rates for extreme mortality securitisations is based on separate modelling of frequency and severity of epidemics. The parameters of the distributions could be based entirely on historical data or be adjusted to reflect current forecasts for both frequency and severity. Binomial distribution would generally be used for modelling frequency. There are a number of approaches to modelling severity. One of them, which was used in analysing the Tartan Capital securitisation, involves modelling epidemic event severity as a percentage of excess mortality fitted to several historical data points. Figure 11.12 shows excess mortality for the US fitted to six severity data points, one of which has been adjusted by placing a cap on broad longevity improvement in the general population.

While the current focus is on modelling flu pandemics, there is an obvious risk of emergence of other diseases. This risk has not been modelled, and possibly cannot be modelled adequately. Its existence, however, needs to be taken into account in pricing extreme mortality bonds and similar securities.

Figure 11.12 Fitted severity curve for excess mortality resulted from epidemics



Sources: Scottish Re and Milliman, Inc.

Probabilistic modelling of extreme mortality in the securitisation context

While probabilistic modelling of the degree of sophistication demonstrated in Figure 11.11 is currently not possible in extreme mortality securitisation, simulation models of the same general type are being developed and accepted as valuable tools in the analysis of extreme mortality. The Vita IV transaction done by Swiss Re at the end of 2009, in part to obtain protection against the H1N1 virus, was structured using the Infectious Disease Model developed by Risk Management Solutions (RMS), the firm known primarily for its expertise in modelling property catastrophe risks.

The RMS model also incorporated the probabilistic analysis of the H1N1 pandemic, including possible mutations and antiviral-drug-resistance scenarios. The use of a probabilistic model of the type developed by RMS has the potential to grow the market if it gives investors extra confidence in the modelling results for extreme mortality bonds.

Analysis of modelling results

Results of mortality index modelling for securitisations such as Tartan Capital have clearly demonstrated that the pandemic component is the key driver of mortality jumps, accounting for about 95% of simulated losses to investors. If the trigger values are moved closer to the mean – the mortality jumps leading to bond default become less “extreme” – other components could start playing a greater role. There is some indication that this might happen and we will see mortality bonds with higher probability of default.

It has been suggested that extreme value theory (EVT) could be used as an additional tool in quantifying the risk of extreme mortality bonds. It would appear that the EVT approach would be most useful in the very tail of the extreme events distribution, with the trigger point set to the index values with very low probability of occurrence. Even there, the EVT approach cannot replace direct simulations and can only provide a check on simulation results and additional insight into bond pricing.

Scenario testing

Scenario testing is important both for validating the models and for determining sensitivity of results to changes in some of the parameters. In modelling extreme mortality bonds, one would be remiss in not ascertaining the effect on the bonds of such events as a repeat of the 1918 flu pandemic.

While a probabilistic framework is inherently better than any deterministic analysis, scenario testing could add significant value in analysing extreme mortality bonds. Currently available stochastic models are based on

numerous assumptions and, out of necessity, utilise very limited historical data. Carefully chosen scenarios provide a check on the results of stochastic modelling as well as another way to assess the risk of extreme mortality. Choosing scenarios to perform stress testing is also important in providing the full picture of the risk involved in investing in extreme mortality bonds.

MORTALITY DERIVATIVES

A number of derivative instruments could be based on mortality risk. These instruments are not limited to the derivatives linked to extreme mortality and could cover even relatively small changes in mortality levels. In addition, they could have a tenor much longer than the three years typical for extreme mortality bonds. The longer time horizon permits the transfer of true longevity risk. (Longevity risk transfer is covered in more detail in Chapter 15.)

Mortality swaps present an example of a derivative based on mortality experience. In a mortality swap contract, counterparties swap a predetermined payment or series of payments for payments whose amounts are based on the number of deaths/survivors in a given cohort. There have already been some private mortality swap transactions.

Another example of a mortality derivative is mortality options. In these contracts, the payout is a function of the mortality index value on a given date. The key to the growth of mortality derivatives is establishing liquidity and standard reference populations or indexes.

ADDITIONAL CONSIDERATIONS FOR INVESTORS

Some of the securities with payments linked to extreme event risk could have very low correlation with other capital markets instruments. Property catastrophe bonds (described in Chapter 3) are a good example of such securities, having relatively weak correlation with financial markets. Extreme mortality bonds and other securities linked to big jumps in mortality are not in this category, however. A true extreme mortality event such as a pandemic could lead to economic and social disruption that would affect all financial instruments. In such an extreme event, most risks suddenly become correlated. This limits the diversification benefit of introducing extreme mortality linked securities into an overall investment portfolio. The lower diversification benefit, compared with property cat bonds, should be reflected in the price investors pay for extreme mortality instruments.

If the trigger point is set lower, events of lower severity now able to

trigger the bond will likely not cause the kind of turmoil in financial markets that would result from a pandemic. In this case, there is weaker correlation of these securities with other financial instruments, and greater benefits to investors to be obtained from diversification.

TRENDS AND EXPECTATIONS

We have seen only a small number of transactions transferring true mortality risk from the insurance industry to the capital markets. However, their number is going to grow. Insurance and reinsurance companies are becoming aware of this method of risk transfer and its advantages. Investors, as they are learning how to analyse securities based on the risk of extreme mortality, are becoming more interested in investing in this asset class. The key reasons for the anticipated growth in extreme mortality and longevity securitisations are as follows.

- ❑ There is a growing realisation that the risk of extreme mortality is real and probably increasing. The implementation of the enterprise risk management approach throughout the life insurance industry brings additional attention to the magnitude of this risk. Transferring some of the risk of extreme mortality to investors is a natural choice for the life insurance industry. Additional scrutiny on the part of rating agencies and regulators provides further impetus for securitising this risk.
- ❑ With the first mortality bonds issued and the most difficult structuring issues resolved, it will be easier for more of these securities to be issued in the future.
- ❑ Investors are becoming more comfortable with, and better educated about, the risk of extreme mortality. The relatively high returns offered by extreme mortality bonds serve as an attractor for investors in their ubiquitous search for alpha.

Other important developments that will affect the future of the market are as follows.

- ❑ With the number of extreme-mortality-linked securities growing, a secondary trading market is developing, providing some liquidity to investors.
- ❑ The insurance derivative market is expected to grow, particularly if traded contracts appear in the marketplace.
- ❑ Innovation is expected to continue, especially in the areas of developing

better indexes and constructing new derivative products linked to the risk of extreme mortality. It is possible that we will see exchange-traded mortality securities in the near future.

- ❑ There is an expectation that we will see mortality based securities with a higher probability of default – transferring less extreme mortality risk, but still providing significant risk transfer. There has already been some movement in this direction.
- ❑ Methods of quantifying mortality risk transfer will be refined, and new approaches will be developed. Better ways to quantify the risk will make extreme mortality linked securities more attractive to investors and contribute to the growth in their issuance. The use of stochastic models such as the one utilised in structuring the Vita IV bond in 2009 is expected to continue and grow.

Extreme mortality securitisations will continue to grow. The transfer of the risk of extreme mortality to the capital markets will benefit both the insurance companies laying off the risk and the capital markets participants investing in these securities.

Life Insurance Settlements

INSURANCE POLICY AS A TRADABLE ASSET

A life settlement is usually defined as sale of the ownership of a life insurance policy or its benefits, or the transfer, assignment or bequest of a life insurance policy or of the benefits of a life insurance policy for a consideration by the owner of the policy when the insured does not have a life-threatening medical condition.

A life insurance policy could have value in and of itself beyond providing a payment to beneficiaries in the case of the death of the policyholder. This value exists even if the original purpose of buying the policy is no longer valid and the policy is not needed for its death benefits. A way to realise this value is to sell the rights to the death benefits to another party. If the price offered to the policyholder for a life insurance policy is greater than the cash surrender value of the policy, under certain circumstances it could be in the policyholder's best interest to sell the policy. For an investor, in a simplified view the transaction could make sense if the net present value of the expected cashflows – including the price paid for the policy, future premium payments and the policy benefit – is positive. In other words, a life insurance policy could be treated as a security. There is a long-standing dispute, at both federal and state levels in the US, and also in other countries, over whether an insurance policy should be considered a security from the legal point of view, but from the finance point of view it is one.

The right of policyholders to sell their life insurance policies has been repeatedly challenged in recent years, and there have been numerous attempts to put significant restrictions on such sales. While certain restrictions remain and others might be imposed, the fundamental view of a life insurance policy as the property of its owner, who has the right to sell it, has been firmly established in the US. In fact, some see the issue as having been confirmed a century ago, in the 1911 *Grigsby v. Russell* decision, in which the US Supreme Court stated that "Life insurance has become in our days one

of the best recognized forms of investment and self-compelled saving.” To address a possible objection, the Court further stated, “But when the question arises upon an assignment, it is assumed that the objection to the insurance as a wager is out of the case”; and, further, “So far as reasonable safety permits, it is desirable to give to life policies the ordinary characteristics of property ... To deny the right to sell except to persons having such an [insurable] interest is to diminish appreciably the value of the contract in the owner’s hands.”

While the fundamental right of individuals in the US to sell their life insurance policies has generally been established, there could still be many legal and regulatory issues to be resolved to exercise this right fully. In addition, to exercise the right to sell a life insurance policy, the policy has to be valid. This seemingly obvious point becomes significant in the context of investing in life insurance policies, with the question of validity being tied to that of the insurable interest at the time of issue. This subject is covered later in the chapter.

The discussion of tradable life insurance policies in this chapter includes a number of topics that appear to be irrelevant to investors and of more interest to other participants in the market. It will become clear why even the details of how insurance policies were purchased, possibly years before investors buy these policies, are critical to the assessment of investment risk and valuation of these securities.

LIFE SETTLEMENTS

Life settlements are financial transactions involving the sale of a life insurance policy by its owner to a third party. The buyer becomes the owner of the policy in the sense of being its beneficiary and assuming the responsibility for paying premiums.

For an insurance policy to have financial value to investors, the insured party does not necessarily have to have experienced a significant deterioration in health. For example, many life insurance policies are structured in a way whereby the premium payments remain level even though the rate of mortality increases over time. Effectively, in the beginning the premiums paid are higher than necessary for the expected level of claims. After a certain period, however, the situation reverses and the premiums no longer cover claims and other expenses as mortality goes up with age. The policy is still profitable to the life insurance company because the “overpayment” in the beginning more than offsets the “underpayment” towards the end of the policy term. Reserves that have been built up from the beginning are

used to pay for claims, most of which come later. This simplified example further demonstrates how an insurance policy could have monetary value to the policyholder who has been paying premiums for several years. On the expected basis, the net present value of the future premiums could be lower than the net present value of the death benefit, often by a significant amount. This difference is even greater for a policyholder whose health condition has significantly deteriorated since the initial underwriting, and whose mortality rate has thus increased beyond the expected value. The value of such policies to potential investors has correspondingly gone up.

Evolution of the market

In the 1990s, a significant number of AIDS-afflicted men were in a position where they needed financial resources either to pay for their medical care or to improve the quality of what at the time was considered to be the very end of their lives. Some of them had life insurance policies that would pay upon their death but would not provide any real help when they most needed it. The appearance of investors willing to provide immediate cash in return for later receiving a greater payout from life insurance companies created a market for such life insurance policies. That was the beginning of the era of viatical settlements.

The landscape has changed dramatically since then, and now, many years later, we have a market for life insurance policies that does not involve terminally ill policyholders seeking to cash in on their policies. Many of the policyholders selling or attempting to sell their policies are not sick at all, and their motivation for entering into a life settlement transaction is completely different from that of the policyholders in viatical settlements years ago. The purchasers of the policies have changed as well. The current investor base in life settlements is primarily institutional, with some of the well-known banks and pension funds playing an active role in the transactions.

Life settlements vs. viatical settlements

Life settlements are traditionally defined as separate and distinct from viaticals, and many professionals in the industry take special care to differentiate themselves from those dealing in viatical settlements. Life settlements are defined as the purchase of life insurance policies from policyholders who are not terminally ill even if they are in their old age and sick. It is difficult to draw a bright line between the two categories, but in most cases, if the life expectancy of an insurance policy seller is less than 24 months, the transaction will be termed a viatical settlement.

Table 12.1 Difference between typical viatical and life settlements

	Viatical settlements	Life settlements
Life expectancy	< 24 months	> 24 months Average 5–7 years
Policy face value	< US\$250K Average US\$100K or less	> US\$250K Average over US\$1 million
Health impairments	Terminal stages of AIDS or cancer	Chronic diseases; in some cases health impairments not greater than average for older ages

There is no exact demarcation between viaticals and life settlements; a number of factors in addition to those listed in Table 12.1 can play a role, but the life expectancy is the primary differentiator.

It is important to point out that from the legal point of view the definitions of viatical and other life settlements usually differ. Insurance laws and regulations vary by state in the US, and there are currently some states that do not distinguish between these two categories at all while others provide distinctly different definitions. The definition affects the legal requirements that have to be satisfied when entering into such a transaction.

LIFE SETTLEMENT SECURITISATIONS

The standard securitisation approach of assembling a pool of securities and then slicing it into pieces to sell to investors works for life settlements too. Portfolios of life settlements and even viaticals have been securitised, albeit on a small scale. The one large securitisation that was supposed to pave the way to growth of the market, that of Coventry First/Ritchie Capital Management, was abandoned at the very last moment, after receiving an indicative rating from a leading rating agency, for reasons that seem to have little to do with the general merits of securitising life insurance settlements. The reported 2009-rated private securitisation by AIG of its book of life settlements with the aggregate face value of US\$8.4 billion and netting over US\$2 billion to go towards possible repayment of the government loan could serve as an important catalyst of growth for future securitisations. The ability to securitise large pools of life settlements would lead to the growth of the life settlement market as a whole, but significant obstacles still remain.

LEGAL AND ETHICAL ISSUES

In certain jurisdictions – in particular in some countries in Europe, where there is a large investor base for life settlements – there are still open questions as to whether purchasing a life insurance policy or a fraction of a policy is equivalent to purchasing a security. The answer to these questions affects the regulatory treatment of the transactions and could have an impact on the regulatory and licensing requirements imposed on the funds investing in life settlements. There exists some level of uncertainty even in the US, where these types of transactions typically originate.

Legal and ethical issues surrounding life settlements are of disproportionate importance and have affected the way the market has developed and the types of investors who have become its active participants. Some of these issues continue to affect the investment risk of these securities.

Ethical considerations

When the idea of selling insurance policies to investors was first introduced, there was some concern that policyholders could be taken advantage of by unscrupulous operators. Viatical settlements are undoubtedly an area of potential abuse, which explains why it is tightly regulated in many states. The public view of viatical settlements has always been mixed even when no laws or regulations are violated. Some see viatical providers as performing an important public service by enabling sick policyholders to obtain financial funds when they are most needed, in order to pay for better medical care or simply enjoy their last days. It is seen as a cruel irony that some get access to the money in the insurance policy only in the grave, when it is no longer needed. The ability to monetise the financial value of life insurance policies has indeed helped many people. On the other hand, extreme care should be taken to prevent unprincipled advisers from taking advantage of the sick by persuading them to sell their life insurance policies for a price that is too low, or in situations where the sale of the policy is not in the best interest of the policyholder. Differentiating life settlements from viaticals is important to the life settlement industry that is trying to avoid any appearance of taking advantage of sick people. The difference between the two is real and not limited to semantics.

It has been pointed out that investors have a financial interest in seeing the people from whom they have purchased life insurance policies die sooner rather than later. While nobody would suggest the possibility of an investor committing murder in order to receive the insurance payout, some investors have felt moral reservations that have prevented them from

participating in this business. Such feelings might be justified, in particular in view of some of the abuses that have occurred, mostly in the early stages of the market. It is worth pointing out, however, that this situation is far from unique. For example, providers of some types of annuity products could be seen as benefiting financially from the early death of the annuitants. However, nobody would question the benefits to the society and individuals stemming from the existence of annuity products.

In fact, the opposite argument could be easily made. Let's consider an individual who has a life insurance policy with a large face value. The policy was initially purchased to provide for his spouse in case of his premature death, but the spouse has since passed away. The individual does not have family members who would need financial support in case of his death. He has limited assets but is now faced with significant medical bills. One could easily argue that a financial adviser to such a person has a duty to consider and possibly recommend the use of an asset such as the life insurance policy to pay the medical bills or simply to use the proceeds of the sale to improve the quality of the person's life. A financial adviser not mentioning such an option to his client could even be seen as committing malpractice, in particular if the valuable but no-longer-needed life insurance policy is allowed to lapse or is settled for the small cash surrender value offered by the insurance company.

As touched on above, life settlements are not the only product whose provider can be seen as having a financial interest in an early demise of certain individuals. Many annuity products provide payments to annuitants as long as they are alive, and have any obligations terminated upon death. While it could be said that insurance companies providing these annuity products would generate greater products were the annuitants to die early, it is generally accepted that the annuity products serve an important financial function, and there are no valid ethical objections to them. Annuities providing a predictable stream of payments are an important retirement-planning tool that affords a degree of security to purchasers of these products. A similar case is that of a pension plan that provides defined benefits to participants as long as they remain alive. Pension plans play an important role in the society; there are no ethical concerns or issues involved. The pension plan argument, however, is weaker than that of annuities because pension plans are typically governed by trustees having no personal financial interest.

Ethical, as well as legal, considerations have to do also with protecting the personal data of individuals considering life settlements or having already settled their insurance policies. Detailed personal information is disclosed in

the life settlement process and there is every expectation that the information will remain confidential and not be disseminated. Investors in life settlements have access to this information; they would be in breach of ethical and often regulatory rules if the information is not properly safeguarded.

The individuals considering life settlement transactions tend to be older, retired and in poor health. This vulnerable population has to be protected against the potential of predatory sales practices and unfair pricing. Proper regulation can address these concerns and alleviate objections on the part of socially responsible investors.

There appears to be a growing consensus that life settlements do benefit the society and that policyholders have the right to dispose of their policies in any way they see fit. At the same time, it is undeniable that the life settlements arena should be subject to close regulation to prevent any abuses. While there are obvious ethical issues involved, the natural way to address any concerns is by having a robust regulatory framework governing the life settlement marketplace. Such a framework will also protect the interests of investors by establishing clear rules and reducing the uncertainty.

MARKET PARTICIPANTS

The process of selling or buying a life insurance policy has several steps and has to go through several intermediaries before it reaches the investor. While the terminology is not always consistent, the key participants are described in Table 12.2.

Figure 12.1 illustrates the traditional process flow in a life settlement transaction. The same party can perform more than one function. For example, life settlement broker and life settlement provider might be the same entity. There is a growing trend towards vertical integration. While economically advantageous, such integration has a potential for creating a conflict of interest.

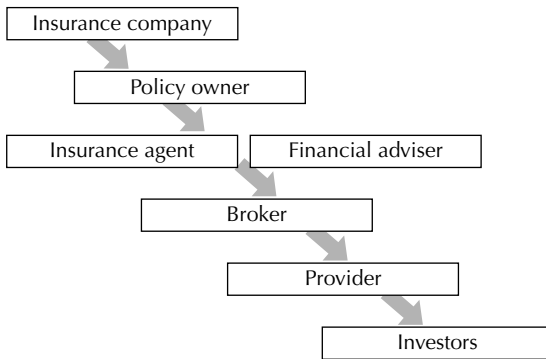
The number of steps and parties involved in a life settlement transaction partly explains why the commissions and fees constitute such a sizable percentage of the total amount paid by investors. The growing transaction transparency is expected to lead to lower payments to third party intermediaries, making the market more efficient and facilitating its growth.

CURRENT AND FUTURE MARKET SIZE

The exact size of the life settlement market is unknown and published estimates have varied widely. The reason for the uncertainty as to the market size is the private nature of life settlement transactions.

Table 12.2 Market participants

Insurance company	The original issuer of the life insurance policy. Must be notified of the transfer of ownership/beneficiary change. Receives premium payments over the life of the policy and pays claims. In some cases serves as investor in life settlements for policies usually issued by other carriers.
Policyholder/insured	Seller of the beneficiary rights to a life insurance policy. Receives a lump sum payment and/or another consideration in return for the right to receive the policy benefit from the life insurance company.
Financial adviser and/or insurance agent	Provider of advisory services and facilitator of the transaction. Could be compensated on a fee or commission basis.
Life settlement broker	Broker facilitating the life settlement transaction. Typically paid a commission for the services. Subject to licensing requirements in most states in the US.
Life expectancy provider (LE provider)/Medical underwriter	Provides review of the medical condition of the insured and associated mortality profile to develop a view of the expected mortality, somewhat similar to the underwriting process of life insurance companies.
Life settlement provider	Purchaser of life insurance policies for investors or for its own account. Typically a separate company or a bank. Makes a payment to the seller of an amount in excess of the cash surrender value of the policy. Typically subject to licensing requirements in the state of residence of the policy owner in the US. One of the parties responsible for addressing compliance issues.
Servicing and tracking agent	Monitors the status and whereabouts of the insured using methods similar to those utilised in the servicing of consumer loans. Provides the information to investors. Could be performing such servicing functions as claim processing and premium payment.
Trust administrator	Responsible for the administration of the trust if one is established for life insurance policies.
Investor or “funder”	Funding source for the purchase of life insurance policies. Could purchase policies from original policyholders in the secondary markets or from other investors in the so-called tertiary market. Typically, an institutional investor such as a hedge fund.

Figure 12.1 Life settlement flow diagram

While investors in life settlements have come from the US, Europe, Asia, Latin America and Australia, the policies they have invested in have almost all originated in the US. The right of policyholders to sell and of investors to buy insurance policies is deemed established in the US but not in most other countries. In addition, the types of life insurance products offered and the pricing dynamics make the US life settlement market particularly attractive to investors. Consequently, estimates of the market potential and the general discussion of the market tend to focus on the US. We do note that there are other markets in addition to the US, such as the market for German traded endowment policies (TEPs). Investors buying German TEPs from policyholders and holding them to maturity receive the terminal bonus payments structured into these products.

The current market size estimates vary and reliable data is impossible to obtain due to the nature of the market. In 2009, the settled amount in force is above US\$20 billion in face value, with some estimates going as high as US\$55 billion. Annual aggregate face value of settled policies grew steadily to exceed US\$10 billion per year until it met a slowdown driven by the 2008–09 financial crisis and the resulting shortage of investment capacity. The market started its slow recovery in 2009 as investment capital began to flow back into the industry.

Projections of the market size in several years reach US\$160 billion and even higher. There are also estimates that expect the market size to be much smaller. However, the consensus opinion is that the market will continue to grow as the product becomes better known and the baby boomers age. Currently, most policyholders are unaware of the option to sell unneeded or unaffordable life insurance policies and many let their policies lapse without

exploring the life settlement option. This is gradually changing as the knowledge of life settlements spreads. Baby boomers growing older also serves to expand the pool of policies that are potential candidates for life settlements. Finally, new investors are exploring the market due to its potential to deliver high risk-adjusted returns and provide diversification to their investment portfolios.

REGULATORY ISSUES

Laws and regulations governing life settlements are still evolving. Relevant regulations in the US differ greatly from state to state: they are very comprehensive in selected states and largely nonexistent in others. There are still some states where regulations lump life settlements together with viaticals. In others, the primary regulations are those designed for consumer protection without specific reference to life settlements. Most states have life settlement regulations either adopted or pending in the state legislatures. Unfortunately, these new regulations are not uniform either and differ from state to state, with two primary variations being the most common. The regulations correctly focus on protecting individual consumers; only in isolated cases is attention paid to providing protection to investors in life settlements.

While the urgency to enact life settlement regulations is driven primarily by the potential for abuse and the need to protect consumers, having a clear regulatory framework will be beneficial for all parties involved in life settlement transactions, from individual policyholders to insurance companies to investors. It is important to keep in mind, however, that hastily enacted legislation could have unintended consequences: regulation aimed at preventing specific abuses in life settlements has the potential of making it harder to engage in legitimate life settlement transactions. This would hurt consumers by depriving them of a valuable financial option.

In addition to insurance regulations that vary by state, there are also securities regulations that might be applicable to life settlement transactions. While the scope of the regulations is unclear and subject to an ongoing debate, settlement of a variable insurance policy is considered to be a securities transaction by the Financial Industry Regulatory Authority (FINRA), since a variable life insurance policy is treated as a security. As such, those facilitating a settlement of a variable life insurance policy are subject to specific obligations regarding due diligence, suitability, execution and compensation. This includes the case of trading already settled policies between investors.

THE LINK BETWEEN INVESTOR RISK AND CONSUMER PROTECTION

While investors, like everybody else, want to see consumer interests protected, in the case of life settlements assuring that consumer protection laws and regulations have been followed is also a critical part of the due diligence process. Any impropriety or appearance of impropriety could endanger the value of a life settlement investment. This type of due diligence is particularly important because there is significant confusion as to the interpretation of the existing regulations and vast differences in regulatory treatment from one jurisdiction to another. The regulatory environment is still evolving for this new product. One of the consequences is that best practices have not been fully established and are continuing to evolve as well.

The two points in time that warrant particular attention in the relevant part of the due diligence process are the purchase of the life insurance policy by the policyholder from the insurance company and the point of sale by the policyholder to the investor, the latter usually being the point when the investor due diligence process takes place. The slightly more complicated situation when a policy is being traded between investors is discussed in subsequent chapters as part of the overview of the tertiary markets.

The examination of how the policy was originally purchased by the policyholder is focused primarily on whether the policy is a so-called STOLI, which stands for stranger-originated life insurance policy. The STOLI issues are discussed later in this chapter and involve, among other things, potential for the consumer to be taken advantage of by a broker who would like to generate commissions on the sale of the policy, and might suggest to a consumer uneducated in financial and legal matters the idea of purchasing a life insurance policy and then flipping it to investors, either immediately or when the contestability period ends. This raises the question of the insurable interest at the time of issuance, and likely invalidates the insurance contract. If this happens and the policy is judged to be invalid after it has been settled, the investor might find himself in an unenviable position of trying to prove the validity of the policy in court, or trying to recover some of the losses from the intermediaries or the insured. The investor also risks finding himself involved in litigation initiated by the policyholder against advisers and brokers. This could be the case even if the policy is not judged to be STOLI. The consumer might not have been fully informed of the tax consequences of settling a life insurance policy or of the reduced ability to buy additional insurance in the future. Since individuals settling their policies tend to be elderly and in frail health, there is the potential of diminished capacity at the time of making the decisions and signing the documents.

This situation also increases the risk of a lawsuit by the survivors of the insured, who might have been unaware that the policy had been settled and be expecting to receive death benefits instead of the benefits going to investors. In this case, the questions again would concern informed consent and the capacity of the insured to enter into a life settlement transaction. To illustrate the risk this presents to investors, it is enough to bring up the example of a policyholder in the early stages of Alzheimer's disease settling the policy before being diagnosed. An investor and every party to the transaction should also consider the possibility of claims that the policyholder did not understand the full implication of the transaction even when the policyholder is well educated and financially savvy. A good example of the risk is the case of the CNN television broadcaster Larry King, who allegedly was persuaded to engage in life settlement and possible STOLI transactions without fully understanding their implications.

Inadvertent disclosure of consumer personal information is another investment risk to consider. It is possible for the risk to be minimised by the performance of due diligence at the time of the original transaction and by the establishment of adequate controls so that personal information will not be accidentally released at some later point.

Investor due diligence prior to entering into a life settlement transaction, when it comes to consumer protection issues, is often focused on the due diligence on the brokers and providers originating the transaction. Having understood the way they operate and the procedures they use, the investor gains greater comfort and more easily relies on them when there is constant flow of new life settlements from the same sources. It has given investors a certain degree of reassurance that their portfolios are homogeneous and have been assembled of relatively similar insurance policies from a few well-vetted sources. While this is a sound approach, one should be aware that, if the reliance on a source of policies and the trust have been misplaced, there is a risk of a large part of the portfolio suffering losses as opposed to the losses being limited to at most a few individual policies.

TAX ISSUES

While life insurance benefits are usually not subject to federal income tax in the US, this exemption does not apply to investors owning a life insurance policy they have acquired in a life settlement transaction. Considering an example of a level premium term life insurance policy without cash surrender value, the Internal Revenue Service (IRS) issued an opinion in May 2009 that investors receiving death benefits on a policy should be taxed

on the proceeds at the ordinary income rates. However, investors could include in their cost basis for tax purposes both the amount paid to acquire the policy and other amounts such as premiums paid to the insurance company after the policy was settled. In an unusual but possible situation when the net income is negative – that is, the policy benefit is lower than the sum of the amount paid for the contract at settlement and the premiums paid for it subsequently – no tax is due.

If the investor decides to sell the policy before maturity in the tertiary markets to another investor, the gain realised is treated under capital gains rules with the cost basis calculated the same way as when the policy is held to maturity.

The same IRS opinion clarifies that non-US-based investors are subject to the same tax – since the income is derived from sources within the US – and, consequently, are also subject to US tax withholding on the proceeds unless they are domiciled in a jurisdiction that has a tax treaty with the US. This affects the funds domiciled in offshore or other tax-advantageous jurisdictions. Legal and tax advice should be obtained by investors before engaging in a life settlement transaction since there are many intricacies and emerging issues in the legal and tax treatment of these instruments.

The insured settling a policy pays taxes on the settlement amount that are treated as capital gain after certain adjustments are made to the cost basis. In the case of a policy with a cash surrender value, this value is considered to be ordinary income, while the excess of the cash surrender value is considered capital gains, all after properly calculating the tax cost basis. Legal and tax advice should be sought by the policyholder in any specific transaction. As with other aspects of life settlements, the investor is advised to perform some due diligence on the other parties involved in the transaction to make sure that laws and regulations have been followed, including the advisers making the insured aware of any tax implications of life settlements. Otherwise, investors might end up with the risk of a lawsuit resulting in investment losses and reputational damage.

INSURABLE INTEREST

A critical issue to be considered in a life settlement transaction is that of insurable interest. The investor has to make sure that insurable interest existed at the time the policy was issued and that the policy is not a stranger-(or investor-) originated life insurance. Otherwise, the insurance contract might be unenforceable, resulting in investment loss.

Insurable interest is a key concept in insurance contract law. It is

supposed to serve the dual purpose of defining insurance in legal terms and preventing certain activities detrimental to the welfare of the society. Insurable interest is a complicated concept, the exact definition of which depends on the jurisdiction. For example, the California Insurance Code, Section 10110, defines insurable interest in the following way:

An insurable interest, with reference to life and disability insurance, is an interest based upon a reasonable expectation of pecuniary advantage through the continued life, health or bodily safety of another person and consequent loss by reason of that person's death or disability or a substantial interest engendered by love and affection in the case of individuals closely related by blood or law.

An individual has an unlimited insurable interest in his or her life, health or bodily safety and may lawfully take out a policy of insurance on his or her own life, health or bodily safety and have the policy made payable to whomsoever he or she pleases, regardless of whether the beneficiary designated has an insurable interest.

Additional examples are further provided. It is also stated:

Every person has an insurable interest in the life and health of:

- (a) Himself.
- (b) Any person on whom he depends wholly or in part for education or support.
- (c) Any person under a legal obligation to him for the payment of money or respecting property or services, of which death or illness might delay or prevent the performance.
- (d) Any person upon whose life any estate or interest vested in him depends.

Insurable interest is needed for a life insurance contract to be valid. The same section of the California Insurance Code states:

Any contract of life or disability insurance procured or caused to be procured upon another individual is void unless the person applying for the insurance has an insurable interest in the individual insured at the time of the application.

The law requires that insurable interest exist at the time of application for and issuance of a life insurance policy. This is different from what is required for indemnity contracts, such as property and casualty insurance, where insurable interest should also exist at the time of the loss.

The two reasons why insurable interest is required by law in life insurance are to avoid the moral hazard of owning a policy on the life of a stranger, and to prevent wagering. While the idea of an investor in a policy murdering the insured appears somewhat ridiculous in modern society, it is

a good reason to ensure the existence of insurable interest. The notorious case of the Blue-Eyed Six in 19th Century Pennsylvania involved the purchase of a life insurance policy on the life of another person where no insurable interest existed. The insured was murdered to collect the policy benefit, and five out of the six accused were convicted and hanged. Avoiding gambling or wagering by investors on the life of a stranger is also against public policy and constitutes the second reason why the concept of insurable interest has been part of the law for centuries.

To provide a more simple and practical definition, we can say that insurable interest exists in the following four categories of cases:

1. Relations by blood or marriage

Individuals are presumed to have insurable interest in the lives of their spouses and dependents and to be interested in their welfare. This generally includes husbands and wives, parents and children, brothers and sisters, grandparents and grandchildren. In most cases, this does not include cousins, uncles and aunts, nieces and nephews, stepparents and stepchildren, or relatives by marriage as opposed to by blood. Of course, a person is also presumed to have insurable interest in his or her own life.

2. Business relationships

Business relationships could create a financial interest in the continuing welfare of an individual. For example, a corporation could have insurable interest in the life of its officer or employee, and could buy a key-person life insurance policy on the employee. Similarly, a partner in a business partnership could purchase a life insurance policy on the life of another partner; the policy could also be purchased directly by the partnership.

3. Creditor – debtor relationship

Creditors are presumed to have insurable interest in the lives of their debtors, with the insurable interest being capped at the value of the loan. Sometimes, a debtor's agreement might be required for the purchase of life insurance.

4. Other relations of direct financial dependence

Under certain circumstances, a person who is, to some degree, financially dependent on another person would have insurable interest in the life of that person. An example important to the field of life settlements is the relationship between a charitable organisation and a donor, where the existence of insurable interest is controlled by the specifics of the relationship and the jurisdiction.

INVESTOR- OR STRANGER-ORIGINATED LIFE INSURANCE POLICIES

Stranger-originated life insurance policy, or STOLI, is an arrangement whereby investors, directly or through third parties, encourage individuals to purchase insurance on their lives with the intention of selling the policies to investors for profit. The profit could be direct monetary compensation or some amount of life insurance coverage that is “free”. When defined this way, STOLI is unlikely to meet the insurable interest requirements. Numerous other acronyms have developed, with most of them carrying largely the same meaning. These are speculator-initiated life insurance (SILI), investor-initiated life insurance (IILI), stranger-owned or -originated life insurance (SOLI), and investor-owned or -originated life insurance (IOLI), all them generally having the meaning of policies that were originated with the intention to sell them to investors as opposed to traditional life settlements owned by investors.

In a situation when individuals purchase a life insurance policy using their own funds, it is usually impossible to ascertain intent at the time of purchase, since being aware of the option and seriously considering the sale of the insurance policy to investors does not necessarily imply lack of insurable interest. For this reason, it might or might not be STOLI, where STOLI is defined as a transaction clearly lacking insurable interest and thus invalidating the policy.

Premium financing and the STOLI issue

Another typical case is taking a non-recourse loan used to finance premium payment for the first two years of the policy or a slightly longer period of time. Two years is the contestability period, after which there is much lower chance of an insurance company refusing to pay policy benefits due to irregularities at the time the policy was issued. The policy could be placed in a trust to serve as collateral for the loan. If the insured dies before the loan term ends, the loan balance is paid out of the death benefits received from the insurance policy. Otherwise, at the end of the loan term the policyholder has three options. The first is to repay the loan and keep the life insurance policy, which also involves paying future premiums out of pocket. Another option is to settle the policy by selling it to investors. Part of the proceeds from the sale is then used for repaying the loan. A third option arises when the insured does not want to or cannot afford to repay the loan and keep the policy, but selling the policy to investors would not generate funds sufficient for repaying the loan. Depending on how the loan was structured, in a situation like this the

insured usually has the option of giving the policy to investors and thus being released from the obligation to repay the loan.

To understand whether there are questions regarding insurable interest and the validity of the policy, one has to look at the moment the policyholder applied for and was then issued the policy. Sometimes there are indirect indications that the policy is STOLI and that at the time of policy issue there was a clear intent by the insured to sell the policy at the end of the loan period. For instance, the loan could have an interest rate set at a level above market while the policyholder does not seem to have financial means to repay the loan in two years. This might be interpreted as an indirect indication that the original intent was to settle the policy to investors, and that keeping the policy was never considered. Evidence of communication between the insured and the promoter or initiator of the premium transaction that identifies the purpose of the life insurance purchase provides more direct proof, but such evidence is rarely available. Even if the intent was to sell the policy to investors at the end of the loan term, in some cases it is argued that the intent also included using the life insurance policy to protect the insured's life during the first two or three years after the policy was issued; this "mixed" intent is seen by some as evidence that the insurable interest existed and the policy is valid.

Most premium-financed life insurance policies are not STOLI and are perfectly legitimate. STOLI policies damage the industry's reputation, and the leading players in life settlements are active proponents of enacting strict regulations to clearly define and to prohibit such transactions. The uncertainty as to the exact definition of STOLI has created a significant risk to investors and complicates the due diligence process. The so-called carrier-approved premium finance programmes reduce but do not eliminate this risk to investors.

Wet paper

Sometimes life insurance policies are sold to investors immediately after they have been issued. In such cases, investors should be aware of the increased STOLI risk.

Wet paper is less common now than it was in the past, due to unwillingness on the part of most investors to take on the STOLI risk embedded in such policies. It is important to reiterate, however, that the fact of a policy's being settled shortly after issue does not by itself indicate the lack of insurable interest or imply any other irregularities. Due to the current regulatory developments, there is also a good chance that the transfer of policy owner-

ship within the first two or even five years after the policy is issued will be restricted in some jurisdictions, making the wet-paper question moot.

CONTESTABILITY

A life insurance policy contestability period is the time limit after issuance during which the insurance company can dispute the validity of the policy on the basis of mistake or fraud committed in the application process. The period is typically two years, even though there have been some legislative proposals to increase the period to five years, at least in respect of certain types of application fraud or mistake. During the contestability period, a death claim could be denied or the policy rescinded. Depending on the state and specific policy contract language, there are two main types of contestability clause, one with no exception for fraud and the other with fraud exception. The fraud exception states that the policy shall not be contested by the insurance company after the two-year contestability period in the absence of fraud. The more common type of contestability clause states that the policy shall not be contested once it has been in force for two years from its issue and does not include any fraud exemption. Of course, the company always retains the right to contest the policy and deny claims for nonpayment of premiums. Insurance statutes in some states could invalidate fraud exemptions even if they are part of a signed insurance contract.

The issue of contestability is important in life settlements because of the possibility that the insurance company could claim lack of insurable interest and rescind the policy. Transferring policy ownership to investors invites the scrutiny of how the policy was originally purchased and whether at the time of purchase there was intent to sell the policy to investors, making the policy STOLI. In addition to the risk of STOLI, there is also a possibility that the insured misrepresented their medical history or some other important fact to obtain the policy at lower rates, or at all. Most of that risk goes away when the contestability period is over.

Investor interest in policies within their contestability period has been gradually diminishing due to inability to fully identify and quantify the risk or to be compensated for assuming it.

The end of the contestability period does not mean the end of the risk that the policy could be rescinded or death benefits denied. Some of the risk never goes away. An insurance company might deny a claim many years after policy issue because of an alleged lack of insurable interest at the time the policy was issued. The courts and regulators tend to side with consumers and against insurance companies in such cases. However, when

the policy is owned by investors and not an individual, there is a greater chance that in some cases the courts would uphold the claim denial. The case law is still evolving and it is important to be aware of this risk in the investor due diligence process.

TRUST STRUCTURES AND INVESTOR DUE DILIGENCE

Quite often a life insurance policy is placed in a trust or is owned by an entity as opposed to an individual. The policy might be purchased by a trust or another entity or be transferred into a trust in a life settlement transaction.

A policy could be purchased on the life of an individual by a corporation, a limited liability company or a limited partnership. The burden is on the investor to ensure that the entity had insurable interest at the time of policy issuance, that the entity has the requisite authority to sell the policy, that the documentation is appropriate for the type of entity and life settlement transaction, and that the signer is duly authorised to have the entity enter into the transaction. A policy can also be owned by a trust, requiring a similar type of examination. In the case of trust-owned life insurance policies, due diligence might be even more complicated. It is necessary to understand how the trust has been created, what its purpose is and under what circumstances it purchased the policy. Trusts established by charitable organisations present a particular problem. Significant attention has also been paid to the irrevocable life insurance trusts (ILITs). Such a trust is established specifically for the purpose for taking out a life insurance policy on the grantor. Establishing an irrevocable life insurance trust is used mostly as a tax transfer strategy. It is utilised in estate planning to avoid the policy benefits being considered part of the descendant's estate and thus being taxed at full value. While largely the same result could be accomplished by transferring the policy ownership to another party, if the insured dies within a three-year period after the ownership transfer, the policy benefits are still considered to be part of the estate from the tax point of view by the IRS in the US. Establishing an irrevocable life insurance trust to own the policy from the beginning avoids this tax liability if the trust is properly structured and the insured does not exercise substantial control over the trust and does not possess ownership of the policy. ILITs could also address other estate concerns and provide the flexibility needed in case of multiple beneficiaries. Even though such a trust is typically unfunded or only minimally funded and relies on the insured to make gifts to enable premium payments, it provides a sufficient degree of separation to be considered independent from the insured.

The trust might sometimes be allowed and wish to sell the policy to investors. The reasons vary and could include the policy no longer being needed for estate or other purposes, inability to pay premiums or the decision to allocate funds to purposes of greater relative importance than premium payment and policy maintenance. Another scenario is the term of the policy ending and the decision being made that exercising the conversion option and selling the policy would be the optimal outcome as opposed to collecting the cash surrender value, if any, or maintaining the policy after the conversion. The emergence of the life settlement option has brought attention to the question of the existence of insurable interest for such policies in general. Even though irrevocable life insurance trusts are a relatively common estate-planning tool, the separation mentioned above has now raised questions as to whether such a trust owning an insurance policy did have an insurable interest in the life of the grantor when the policy was issued. While the answer to the question is generally positive, it is important to review the relevant statutes and case law in the jurisdiction where the trust is located.

A life insurance policy could be placed in a trust if it is pledged as collateral for a loan, for example in the context of premium financing. As mentioned above, premium financing requires very careful due diligence on the part of the investor, since there is a greater chance of the policy having been purchased with an intent to later sell it to investors – that is, lack of insurable interest at the point the policy was issued. The terms of the premium-financing loan sometimes indicate that the likelihood of the loan being repaid was very low from the very beginning, and the likely intent of both the insured and the loan provider was for the policy ownership to be transferred to the debtor to be subsequently sold to investors. As a rule, discerning intent is usually challenging if not impossible, and many investors, having discovered a potential problem in their due diligence process, would decide not to take the risk even if they think that the insurable interest existed.

Some investors rely, entirely or in part, on life settlement providers to perform the necessary due diligence process. Full reliance on a provider introduces a level of risk that a prudent investor would not want to assume; functions that are critical should be performed in-house and not be outsourced.

THE USE OF NOT-FOR-PROFIT ORGANISATIONS IN LIFE SETTLEMENTS

Many charities and other not-for-profit organisations receive, as part of their fundraising, life insurance policies from donors who either no longer need

the policies or prefer the charities to benefit from them. Changing the beneficiary to a charity is one way to do this; another is to formally transfer the policy ownership as a gift by making an absolute assignment of the policy. Tax deduction is then provided for at least the cash surrender value of the policy and, if the individual continues to pay premiums on behalf of the charity, for the value of these premiums.

The charity receiving the gift has several options available to it. The three main options are the same as in the case of an individual owning a life insurance policy. They are: keeping the policy until maturity by paying the premiums; letting the policy lapse and receiving cash surrender value if any; and settling the policy by selling it to investors. The last option might be particularly attractive if the charity finds it difficult to make premium payments and administer the policy for an indeterminate period of time or if it has an immediate need for funds to support its activities, and if the value of the policy exceeds any cash surrender value it may have. The relevant option is that of settling the life insurance policy, in particular because such policies tend to have higher-than-average face values and consequently be of more interest to investors. Charitable organisations are a significant source of supply of policies to the marketplace. Sometimes a not-for-profit organisation might have a number of life insurance policies that could form a portfolio to be sold as a whole to investors. These transactions are becoming more common; several not-for-profit organisations, including universities, have even made public announcements about their having sold insurance policies to investors.

Under certain circumstances, a charity could take out a policy on a major donor and have insurable interest. Such a policy could be settled as well.

One of the ways charities are involved in the life settlement space has to do with purchasing life insurance on their donors with premiums financed by investors. A number of legal structures are used; some of them are controversial and involve the questions of insurable interest and STOLI. One common structure involves the following steps.

1. An arranger finds and comes to an agreement with a charity willing to insure some of its donors and be part of the transaction. The charity will receive a portion of the death benefit when the insured donors die. Sometimes there is an upfront payment to the charity serving as an inducement.
2. The charity establishes a trust.
3. Investors put capital in the trust. This could be done through a fund leveraged with debt when financial leverage is available.

4. Life insurance policies are taken on the donors by the trust with the permission of the donors. Typically, the face value of the policies is large. Simultaneously, the trust might also take out single-premium immediate annuity contracts on the same insureds as a way to fund future premium payments and, in some cases, pay interest to investors who have provided the initial funding for the trust. Extensive shopping is usually done to obtain the most favourable life insurance rates by finding insurance companies that would put the applicants in the lower risk categories in their underwriting process. If annuities are also bought, there is usually an attempt to take advantage of the arbitrage opportunities that might result from the different pricing assumptions used by insurance companies for life insurance and annuity contracts.
5. When a donor dies, the policy benefit is received by the trust. It is then split between the investors and the charity, with most of it going to investors so they can recoup their original investment and earn a profit. The transaction could be structured in such a way that the investors receive a predetermined payment and the remainder goes to the charity. An alternative, which might involve using a slightly different structure, is the outright sale of the policies to investors if the charity has unanticipated cash needs in the future, which can happen before or after the contestability period ends. Each has its own benefits and risks and requires legal advice.
6. If the life span of the insureds is longer than projected, the investors could earn a lower-than-projected return or lose money. These transactions are rarely structured in such a way that the charity is risking a loss of this kind.

Such an arrangement could work to the advantage of both investors and the charity if structured correctly. If it is not, it could lead to a multitude of problems and potential losses for all the parties involved. The biggest question is whether there is insurable interest in the transaction. If the transaction is structured to benefit primarily the investors and not the charitable beneficiary, an argument could be made that an insurable interest does not exist. The involvement of a charity does not automatically legitimise such a transaction in the eyes of the law or insurance regulators. Any such transaction should be structured very carefully also in order to protect the charity involved. If the charity is seen to be accepting an inducement, it risks losing its tax-exempt status. In one state, insurance regulators specifically opined on a proposal to raise funds through a securitisation involving the purchase

of life insurance and annuity contracts on a pool of individual donors with the income from the annuities and life insurance benefits being first used to pay interest to investors and only a small portion of the value of the policies going to the charity. The opinion was that the transaction would lack insurable interest and be contrary to the law, since it would be structured to benefit primarily the investors and not the charity. This highlights the importance of making sure there is insurable interest and carefully reviewing the whole transaction.

INVESTOR PERSPECTIVE

From the investor point of view, there could several reasons for investing in life settlements. One reason, just as in the more traditional securities such as corporate bonds, is a set of advantageous risk/return characteristics: many investors see in life settlements the promise of high returns relative to the investment risk.

Another potential advantage of life settlements is the low degree of correlation with traditional financial assets. The main determinant of profitability in properly structured life settlement investments is mortality. Mortality, except in the more extreme cases, has a low degree of correlation with traditional financial markets. However, it is important to point out that life settlements as investments are not uncorrelated with the rest of the financial markets. For example, interest rate risk is present in life settlement investing. The only claim that could be made is that they have a lower degree of correlation relative to the more traditional asset classes. This, however, is still a very important claim and rationale for investing in life settlements.

Finally, a reason for investing in life settlements is the ability to generate returns in excess of the level that would be expected in the efficient markets universe. Generating alpha is predicated on the ability of the investor to identify these inefficiencies and mispricing, and to take advantage of them within a properly constructed and executed investment strategy. This final reason is valid only for the sophisticated investor. The complexity of the required analysis is often underestimated, leading to unpleasant surprises to the less sophisticated portfolio managers and investors in their funds.

Investing in life settlements has also been advocated as an asset-liability management tool for pension funds and other institutions. This argument is rather weak, in particular for pension funds whose liabilities are in most cases inversely correlated with life settlement returns. The reason for this is that greater population longevity increases pension fund liabilities while also leading to diminished cashflows from life settlement investments.

Historical investment performance

The asset class is new and the data on investment performance is very limited. Also, this asset class has been evolving and changing to the degree that past performance is not representative of the expected future results. Historical record of investing in life settlements should be treated with caution.

Overall, investment performance has been mixed. If we are to include viatical settlements in this category, it is a well-known fact that many investors suffered significant losses, even though some realised sizable profits. Unlike viaticals, life settlements are longer-dated investments requiring a significant time period to elapse before the ultimate investment returns become known.

It often takes a long time to know with certainty whether a portfolio of life settlement investments is performing as expected. If assumptions made at the point of purchase relating to mortality and other factors are incorrect, it is rarely immediately obvious. When performance is different from expectations based on these assumptions, that by itself is not necessarily a sign of the assumptions being wrong. The issue of evaluating historical performance is closely tied to the issue of pricing policies for the inclusion in the investment portfolio. If the historical performance appears to be satisfactory from the profitability point of view, it supports the validity of the pricing assumptions used in the evaluation of life insurance settlements and the lack of a bias leading to potential underpricing. In other words, there should be a feedback between evaluating actual performance of the existing portfolio and pricing of new policies. In fact, the two are best seen as part of the same process.

Valuing a portfolio of life settlement investments also has other important implications, some of which are discussed below. Modelling portfolios of life settlements and other mortality-linked securities is covered in subsequent chapters.

Portfolio valuation

Determining fair value of life settlement investments has been a challenge for more than one hedge fund trying to determine correct net asset values (NAVs) for reporting to their investors. In fact, it is a challenge for every investor. In the rather illiquid market of life settlements, determining fair value for every policy in an investment portfolio based on its market value is unrealistic, especially when it has to be done quite often, as in the case of the funds reporting their NAVs to investors on a monthly basis. Using the

language typically employed in defining fair value, we can say that it is usually impossible to determine the price that would be received for a life settlement asset in an orderly transaction between market participants at the measurement date. In such cases, mark-to-market usually reverts to mark-to-model. The main alternative approach, that of determining fair value based on the market value of identical or similar securities, is difficult to implement directly because each policy is different, and finding another policy that is even remotely close to being identical and at the same time has a readily ascertainable market value is rarely possible. However, even the mark-to-model approach should, wherever possible, incorporate external inputs such as the available market pricing for other life settlement policies. Leaving aside the broader question of whether market value is always the best estimate of fair value, it is important to note that some subsegments of the life settlement market could be viewed as distressed. Using market inputs in such cases could produce unanticipated results.

In general, the mark-to-model approach to determining fair value has been heavily criticised because of the high level of subjectivity involved in modelling and the potential for the manipulation of results. However, the current life settlement markets, with their illiquidity and few opportunities for price discovery, are forced to rely on the mark-to-model valuation to a very significant degree. Another complicating factor is the lack of established approaches to modelling and the wide range of assumptions used in life settlement portfolio valuation, leading to widely disparate results.

Risk of overstated portfolio values

We are often happier from ignorance than from knowledge.

François de la Rochefoucauld

When valuation models are crude or nonexistent, mark-to-model could easily become mark-to-make-believe. This leads to a lack of confidence on the part of some investors in the credibility of reported NAVs. It also highlights the inefficiency of the market and the low level of sophistication of many participants. There is a strong need to improve the quantitative approaches used in valuing life settlements and to broaden the implementation of the more advanced of the existing approaches. It is also important to incorporate non-quantitative factors such as legal risks in the modelling process.

Portfolios with longer average life expectancies are subject to being misvalued for an extended period, especially in the current environment of limited use of advanced modelling techniques and limited testing of the

assumptions used in pricing and valuation. Many portfolios of life settlements have performed extremely well; others have fared worse. Some portfolios perform poorly while investors in these portfolios remain unaware of that fact. On paper, the performance can look significantly better than in reality. Since the reported and the real results converge only years in the future, the overvaluing of a life settlement portfolio can continue for a long time.

There is a strong need to use robust models in valuing life settlement investment portfolios as well as in pricing. Best practices in portfolio valuation have not yet been established and the current lack of established valuation methods is troubling. Sophisticated models do exist; they are just not being widely used and many portfolio managers are unaware of their existence and of the weaknesses in the models and approaches they are using now. We describe modelling techniques and valuation approaches in subsequent chapters.

The more advanced modelling approaches to portfolio valuation are closely integrated with those used in initial pricing. This consistency is particularly important if the life settlement portfolio is more actively managed as opposed to the buy-and-hold approach most commonly utilised by life settlement investors.

Competitive advantages and disadvantages for investors

The low level of investor sophistication is the result of both the infancy of this market and the types of participants involved in its development. The idea of life settlements did not originate in investment banks even though the market has by now become institutional as the volume of transactions has grown. The inefficiency of the market is expected to continue for several years, providing the more sophisticated investors with a significant competitive advantage and turning this competitive advantage into an alpha-generating engine.

The market remains highly inefficient with mispricing widespread and unrecognised. The high level of market fragmentation contributes to the overall inefficiency. Price levels are often inconsistent across market segments because many market participants have a low degree of understanding of the drivers of profitability and an even lower degree of understanding of the risks involved. This is a very strong but warranted statement that highlights the uniqueness of the current life settlement market. Risk management on the portfolio level is rarely performed on an advanced level, even though there are some very sophisticated investors

who are properly managing their portfolios. What makes this market dangerous for unsophisticated investors is exactly what makes it attractive to those who have the expertise to analyse these investments and properly manage their portfolios. The market is likely to go through an upheaval as investment losses in some portfolios ultimately become apparent and lead to the downfall of some investment managers. On the other hand, those who use advanced modelling tools have a critical competitive advantage that is also likely to be sustainable for a number of years.

INSURANCE INDUSTRY PERSPECTIVE

Some insurance companies have seen the advent and rapid growth of life settlements as a threat to their core life insurance business. The reason is the potential for anti-selection introduced by the purchase of insurance policies by investors. There are two main components to the anti-selection. One has to do with changes in policy-owner behaviour when the policies are owned by investors, notably a reduction in policy lapse rates. Investors are unlikely to accidentally miss premium payments or be unable to pay premiums due to financial difficulties. Lapse rates are a contributor to the life insurance product profitability, at least for such products as level term life insurance. Premium rates are set based on the assumption of a certain level of lapses, and a decrease in lapses relative to the assumed rate could lead to lower profitability. The anti-selection manifests itself in the fact that the policies of policyholders who have suffered greater than average deterioration in their health are the most likely to be settled and experience lower lapse rates. These are the policies that are more likely to result in death claims and corrode the total profitability of the product line to the insurance company.

Another potential behaviour change not taken into account by the traditional pricing assumptions used by life insurance companies is the fact that investors are more likely to take advantage of policy features other than the standard death benefit. For example, most of the settlements involve universal life insurance policies. The level term policies being settled are usually first converted to universal life by exercising the conversion options. An owner of a universal life insurance policy could have the ability to vary premiums or take loans against the policy under predetermined conditions. Depending on the specific policy, policyholder health and the level of prevailing interest rates, among other factors, investors are likely to take full advantage of these additional options to the detriment of the insurance company. This difference is generally not taken into account in the pricing assumptions. Also, the issue of STOLI creates additional complications for

insurance companies, since it adds to the mortality arbitrage already introduced by the life settlements phenomenon. Some insurance companies have taken a proactive role in dealing with the STOLI problem. Applications for life insurance now often include questions intended to detect the intent to sell the policy to investors, either immediately after issue or once the contestability period has passed. Some insurance companies have made it difficult for their captive agents to facilitate life settlement transactions, or have attempted to specifically prohibit them from doing so.

Other insurance companies, however, see life settlements as a positive phenomenon. The option of selling the policy adds value to life insurance products and has the potential to attract new customers to life insurance companies.

There are even insurance companies that actively participate in the life settlements market and purchase life insurance policies to accumulate sizable portfolios. This serves as a strong endorsement of life settlements. Life insurance companies are well positioned to be active players in this market. They have the level of expertise, from life insurance underwriting to actuarial, that is lacked by most of the current investors in life settlements.

The dispute within the life insurance industry continues, with some seeing life settlements as a type of “cannibalisation” of life insurance and others believing that life settlements improve market efficiency and add value to consumers. What has been recognised by most is that life settlements are here to stay, and the only disagreement could be about their scope and how to make sure all laws and regulations are followed. For some life insurance companies, the growth in the life settlement market could also create the need to reprice their products.

RISKS TO INSURERS

To expand on the explanation of why growth in life settlements could be seen as harmful to the insurance industry, several additional points could be made. Life insurance company profitability is highly sensitive to the assumptions underlying pricing of its products. In most cases, life insurance products are sold for a long term without the ability by the insurance company to later change the rates or adjust product optionality for policies already on the books. Examples of important pricing assumptions for life insurance are mortality, lapses, expenses and investment spreads. All of them have a potential to be negatively affected by the growth in life settlements, with a resulting decrease of profitability.

Mortality and lapses

Mortality and lapse assumptions are to be considered together in this context. Life settlements could lead to the division of life insurance policies in force into two segments. One of them would include policyholders whose health has deteriorated more than expected based on the general assumptions. These policyholders have higher mortality rates and are prime candidates for life settlements. The other category includes those whose health remains better than the general assumptions suggest. This segment of the policyholders has lower mortality and higher life expectancy. The mortality assumptions for the two segments taken together might still hold.

The lapse assumptions built into pricing apply to both of the segments. The segment targeted by the life settlement industry could see its lapse levels decrease because some policyholders would sell their policies to investors instead of allowing them to lapse. Investors are not going to let a policy lapse due to premium nonpayment, as for them it is a valuable financial asset and they are less likely to be financially constrained in paying insurance premiums until the policy benefit is collected. The other segment, containing those whose health is better than the general assumptions suggest, does not experience a noticeable change in lapse rates due to life settlements.

In other words, life settlements could lead to increased persistency of the policyholders with poor health. This results in higher claims for that segment. Since we assume that lapse rates for the other segment are not affected, the lapse rate for the total group of policyholders decreases and the claim level increases. We can see that, even though general mortality assumptions might be correct in the sense of being based on appropriate mortality tables, the change in the lapse rates leads to greater mortality experienced by the policyholders as a group.

In the more sophisticated approach that could be used by insurance companies, pricing assumptions would contain two or more levels of lapse rates applicable to the segments predicted to emerge later on, distinguished by the health condition of the policyholders. The reason for this segmentation is that policyholders who get sicker are often aware of their decreased life expectancy and in most cases are more likely to value their life insurance policies and not let them lapse. While there are some notable exceptions, this logic calls for lower lapse rates assigned to this segment even without an effect of life settlements, and for higher lapse rates to be assigned to the other segment. In this framework, the effect of life settlements is to decrease even further the lapse rates for those who end up sicker than the average. This, in turn, leads to higher claim levels for the life insurance product as a whole.

Insurance companies base their rates on the underwriting performed prior to the inception of the policies. Life settlement companies have the ability to perform their own underwriting based on the more current medical and other relevant data for policyholders several years after the policy inception date and to incorporate the new data into their estimates of individual life expectancy. This ability to differentiate between insurance policies more and less valuable to investors could be seen as a type of anti-selection directed against the insurance company. All of a sudden, the assumptions built into pricing become invalid and the profitability of the product to the insurance company decreases. The current relative size of the life settlement market is too small to have affected the validity of the life insurance industry pricing assumptions; as it grows, the situation may change.

Investment assumptions

While less of a factor, investment assumptions may have to be adjusted as well if there is a sizable growth in the percentage of policies that end up being settled instead of lapsed. Focusing on the asset-liability management aspect, it could be noted that the duration of liabilities is changed due to the life settlement effect. With fewer policies lapsed, the insurance company should also expect to receive more premiums. This increase in revenues has an offset from the cash surrender value not paid for the policies that have been expected to lapse but instead were settled. The net effect might be either positive or negative. The timing of the cashflows is also affected.

Expenses

The expense assumptions could be affected by life settlements in three ways. One is the increase in expenses associated with policy underwriting and issuance. This is already happening, as there is a growing scrutiny of policies with larger face value to ensure they are not being bought with the intent of later selling them to investors. Many insurers have amended their policy applications to include questions on whether the applicant intends to sell the policy. This closer examination aims to identify and reject STOLI applications by ensuring the applicant would have insurable interest in the policy. It has led to an increase in expenses that is likely to continue. The importance of and resources spent on financial underwriting have also increased, since applicants for STOLI policies sometimes apply for a large face amount coverage disproportionate to their income or the financial needs of beneficiaries.

In addition, some insurance companies have been expending resources

on increasing supervision of their agents to prevent STOLI policies from being originated. Distribution sources new to a particular company are more closely examined to avoid the STOLI issue. These new distribution sources might also be more active in the so-called “table shaving” game, often resulting in mortality anti-selection against an insurance company.

The second contributor to the insurance expense levels is extra monitoring of the policies already written for the purpose of identifying, and in some cases trying to prevent, life settlement transactions. When an insurance company becomes aware of a settled life policy, it could review the underwriting file again to make sure that insurable interest existed when the policy was originally issued. Identifying policies that have been settled is not always easy because of the growing use of trusts. Trust arrangements could sometimes mask the transfer of beneficial ownership to investors.

The third contributor to the increase in expenses is the additional cost at the point of claim payments. This is another point at which, for large-face-value policies, an insurance company might want to more closely examine its files to see whether insurable interest existed at inception and whether the transfer of ownership for a settled policy followed all applicable rules and regulations. Here, as in the previous case, there could be significant legal and litigation expenses involved. There have been cases of insurance companies denying or attempting to deny benefits after the death of an insured. These cases are not common; the more likely scenario is that of an insurance company rescinding a policy shortly after issuance due to the lack of insurable interest.

Insurance companies have incurred and will continue to incur some other expenses resulting from the growth in life settlements. These include lobbying costs by insurance companies trying to limit the scope of life settlements. An example is the attempt to increase the contestability period to five years or to disallow the policyholder right to settle a policy within a five-year period after issuance. Reinsurance costs could also rise with the decreased expected profitability of the primary block and the potential of greater fluctuations in profitability.

CONCLUSION

Life settlements are a growing investment class providing unique advantages to investors wishing to diversify their portfolios. Investing in insurance risk has a limited degree of correlation with other types of investments, and life settlements and similar securities provide investors with exposure to one of the types of insurance risk. Of course, pure insurance risk

is not the only risk involved in investing in life settlements, and it is critical to perform comprehensive analysis before investing in these securities.

The market has become fully institutional and continues to develop. It is expected that it will expand as more investor capital enters the life settlement industry and the supply of insurance policies grows.

Life settlements serve an important societal function in providing liquidity to otherwise illiquid life insurance assets. When done properly, life settlements benefit all parties. Certain concerns on the part of life insurance companies and regulators are justified but can be addressed by establishing clear rules governing life settlement transactions and slowly adjusting life insurance pricing assumptions if necessary. Establishing a consistent legal and regulatory framework is an important part of the market evolution and is expected to serve as a catalyst to its future growth. Such a framework will reduce investment risk and permit its better quantification.

At its current stage of development the life settlement market remains highly inefficient and appears to be in dire need of more sophisticated approaches to modelling these securities. Many of these modelling techniques exist but have not been adopted by the majority of investors. The resulting inefficiencies and mispricing make the market perilous to navigate for the less sophisticated investors while creating unique opportunities for those who can turn these inefficiencies into a source of competitive advantage.

Investors in life settlements and similar securities can be subject to significant legal risks of a type that rarely needs to be analysed in the more traditional assets and even other types of insurance-linked securities. Legal due diligence is an integral part of the life settlement investment process; the ability to identify and quantify legal and other risks is essential to proper pricing of life settlements and managing their portfolios. It is a key element of the risk management process, second only to the analysis of mortality characteristics.

Subsequent chapters provide a more technical treatment of the mortality risk involved, pricing approaches and portfolio valuation techniques for life settlements and other insurance-linked securities. We will also analyse synthetic instruments intended to perform similarly to portfolios of life settlements as well as the available hedging tools useful in portfolio management.

Mortality and Longevity Models in Insurance-Linked Securities

MORTALITY AND LONGEVITY

Performance of life insurance and most annuity products is based on mortality of the policyholders. Mortality assumptions are the key ones in pricing these products by life insurance companies. They are also the key determinants of pricing of such products by investors. Any intelligent investment decision related to securities based on mortality risk requires taking a view on mortality underlying these securities. In addition, mortality assumptions are primary inputs in risk management models for investment portfolios that include securities linked to mortality risk, such as life settlements, synthetic mortality securities and longevity swaps. This chapter introduces and explores the fundamental concerns in the modelling of mortality and longevity risk. This is crucial for the analysis of all insurance-linked securities with embedded mortality or longevity risk.

The broad meaning of mortality in this context is the full probabilistic view of the death probabilities of the insureds in the products included in an investment portfolio. The narrower meaning that is often used is the deterministic view of the death probabilities. One of the differences between the two is that the deterministic view assumes that expected death probabilities are known, and the variability of results stems from random fluctuations based exclusively on these probabilities. The stochastic view considers the possibility that the probabilities themselves are random variables that have their own probability distributions. In a simplified framework sometimes employed by investors, the deterministic view considers only the expected cashflows based on the known mortality probabilities and disregards the variability of results due to the stochastic nature of the mortality process; while the probabilistic view looks at the whole range of outcomes assuming that the mortality probabilities are correct. The simplified deterministic framework is a useful tool; problems arise when the analysis stops at the results it produces and neglects to consider a broader stochastic view.

Longevity is defined as the opposite of mortality, that is, the probability distribution of staying alive in a certain period of time or beyond a certain point in time or age. Lower mortality means greater longevity. In general, we speak of longevity or mortality as risks: the term “mortality” is used when greater mortality is considered to be a risk, and the term “longevity” is used when greater longevity presents a financial risk. An example of the former is the greater-than-expected mortality of life insurance policyholders when considered from the point of view of the insurance company; an example of the latter is greater-than-expected longevity when considered from the point of view of a pension fund. Investors could be exposed to either mortality or longevity risk depending on what insurance-linked financial instrument is being utilised.

MORTALITY RATES

Mortality is typically expressed in terms of mortality rates. Mortality rates are death probabilities usually expressed on an annual basis. Death probabilities depend on the age of a person and, with the exception of some younger ages, increase with age. Mortality rates comprise a mortality table. Panel 13.1 introduces some of the concepts helpful in understanding mortality and longevity models and the relevant actuarial terminology.

PANEL 13.1 BASIC CONCEPTS AND FORMULAS USEFUL IN UNDERSTANDING MORTALITY TABLES, SURVIVAL DISTRIBUTIONS, AND MORTALITY AND LONGEVITY MODELS

Actuarial science utilises terminology and notation that are different from those used for similar concepts in related disciplines. Below we introduce some of the key terms useful for those who have not had exposure to mortality tables and mortality modelling.

The probability that a person aged x will die within t years is denoted by ${}_tq_x$; in other words, ${}_tq_x = \Pr [T(x) \leq t]$, where $T(x)$ is time until death of the life aged x . Let us assume that x is expressed in years. We then define ${}_tp_x$ as the complement of ${}_tq_x$, with ${}_tp_x = 1 - {}_tq_x$. In other words, ${}_tp_x$ is the probability that a life aged x will reach the age of $x + t$; or, ${}_tp_x = \Pr [T(x) > t]$. Since these definitions are always applied to the future and not the past, $t > 0$. The notation is useful and allows for many identities to be written in a simple way. For example, for a period less than 1 year ($t < 1$), $q_x = {}_tq_x + {}_tp_{x-1-t}q_{x+t}$.

Both ${}_tq_x$ and ${}_tp_x$ could be defined in terms of the age-at-death function X of a newborn, with its distribution function $F(x) = \Pr [X \leq x]$, and the survival function $S(x) = 1 - F(x)$, which translates into $S(x) = \Pr [X < x]$. In other

words, $S(x)$ is the probability that the person will attain the age of at least x years. $S(x)$ is then equal to ${}_x p_0$. Even though the probabilities are determined for a newborn, in reality we are only interested in conditional probabilities for a person aged x ; the information prior to age x is usually not relevant in this discussion except for the fact that the person has survived to age x . We can then write

$${}_t p_x = \frac{S(x+t)}{S(x)}$$

$${}_t q_x = 1 - \frac{S(x+t)}{S(x)} = \frac{S(x) - S(x+t)}{S(x)}$$

An important special case used in constructing mortality tables occurs when $t = 1$ year. Then we have ${}_1 q_x$, usually denoted as q_x , being the probability of a life aged x dying within one year. p_x , being equal to ${}_1 p_x$, is the probability of a life aged x being alive in one year, that is, attaining the age of $x + 1$. We then can write

$$p_x = \frac{S(x+1)}{S(x)}$$

$$q_x = 1 - \frac{S(x+1)}{S(x)} = \frac{S(x) - S(x+1)}{S(x)}$$

Let us define $f(x)$ as the probability density function of the age-at-death random variable $F(x)$. Then $f(x) = dF(x)/dx$. Here and elsewhere, we are making certain assumptions about the properties of the functions being discussed, such as that $F(x)$ is continuous.

We then define the force of mortality, μ_x , as the probability density function of X at age x , conditional on survival to age x . The force of mortality is then

$$\mu_x = \frac{f(x)}{1 - F(x)}$$

In the more useful terms of survival function, we can express the force of mortality as

$$\mu_x = \frac{S'(x)}{S(x)}$$

The force of mortality is analogous to the hazard function used in reliability engineering to denote instantaneous failure rate of a system or component. It is important to note that force of mortality cannot be negative.

We can define other variables in terms of the force of mortality. ${}_t p_x$ could be defined the following way

$${}_t p_x = e^{-\int_0^t \mu_{x+z} dz}$$

The survival function could be expressed as

$$S(x) = e^{-\int_0^x \mu_z dz}$$

The relationship

$$\int_0^{\lim \text{ age} - x} {}_t p_x \mu_{x+t} dt = 1$$

is interpreted as that of the probability of a person alive at age x dying between that age and the maximum age a human being can live being equal to 100%. This maximum age is sometimes referred to as the limiting age.

The force of mortality, μ_x , is sometimes useful in modelling mortality and longevity. A number of models have been developed, from assuming that the force of mortality remains constant to the Balducci assumption. These approaches have initially been developed for modelling the survival function for fractional ages. Their applicability to modelling mortality and longevity in the context of life-settlement portfolio analysis is limited; however, they could be used in the initial modelling before more sophisticated approaches are utilised.

It is important to note that there could be significant information about an individual aged x , possibly including full medical and other underwriting data. This affects mortality levels, leading, for example, to many different mortality tables.

Complete life expectancy e_x , or complete expectation of life at age x , is defined as the expected value of future lifetime $E[T(x)]$. Then

$$e_x = \int_0^{\lim \text{ age} - x} t {}_t p_x \mu_{x+t} dt$$

This could also be written as

$$e_x = \int_0^{\lim \text{ age} - x} {}_t p_x dt$$

The median future lifetime of an individual at age x , $M(x)$, is determined by the relationship

$$\frac{S[x + M(x)]}{S(x)} = 0.5$$

In the more practical discrete case, the complete life expectancy of a life aged x is calculated as

$$e_x = \sum_{t=1}^{\text{image} - x} {}_t p_x + 0.5$$

If x is expressed in years, we have made an assumption that an individual on average dies in the middle of their year of death. Precision could be increased by making other assumptions regarding mortality changes through a year, as well as by accounting for a fractional starting age x .

Life expectancy expressed as the expected number of complete years of the future lifespan is sometimes called curtate life expectancy. In the formula above, 0.5 accounts for the difference between the complete and curtate life expectancies.

Survival functions described in Panel 13.1 are familiar to fixed-income investors, since they are analogous to the survival functions used to measure bond defaults. While we have mostly looked at continuous survival functions, in practice actuarial mortality functions are usually used in their discrete form, similar to the traditional way of measuring bond defaults; we note that credit quality survival functions could also be constructed in the continuous form.

MORTALITY TABLES

Mortality tables are constructed based on available historical mortality data. The data allows the calculation of mortality rates, which are later modified to account for a number of factors not reflected or fully reflected in the historical experience. Separate mortality tables can be constructed for categories of people who have distinct mortality characteristics such as males and females, smokers and non-smokers, and those who differ by their overall health condition.

Mortality rates comprising a mortality table allow us to calculate life expectancy, which is a concept widely used in life settlement investments and synthetic mortality and longevity securities. However, in life settlements, life expectancy determination typically incorporates additional information used to make adjustments to the mortality rates.

Panel 13.2 introduces some of the basic terms and concepts used in

PANEL 13.2 BASIC PRINCIPLES OF MORTALITY TABLE CONSTRUCTION

Actuarial science utilises terminology and notation that are different from those used for similar concepts in related disciplines. Below we introduce some of the key terms useful for those who have not had exposure to mortality tables and mortality modelling.

A life table typically includes and is based on the number of members of a quasi-cohort surviving at the beginning of sequential time periods. The convention is to have the initial number of individuals, l_0 , to be equal to 100,000 at birth. l_1 is then the number alive at age 1 and l_{100} the number alive at age 100. From the set of l_x , we can then calculate d_x , the number of the members of the cohort dying between the ages x and $x + 1$, or $[x, x + 1)$. $d_x = l_x - l_{x+1}$. The ratio of those dying between the ages of x and $x + 1$ to the number alive at the beginning of the period, d_x / l_x , is interpreted as the previously defined q_x , the probability of a person aged x dying before reaching the age of $x + 1$. A number of other variables, such as life expectancy e_x at age x , could also be calculated based on the same data and presented in a life table.

The term “cohort” usually implies a deterministic view where the group consisting of 100,000 individuals at birth ($l_0 = 100,000$) remains closed to new entrants and members leave the group only through death. In reality, observing a group of newborns over their lifetimes is not how a mortality table is constructed. Instead, a life table is typically based on the mortality rates estimated for a specific population based on a snapshot of data at a single point in time, or rather within a relatively short time period. When we apply the table to a randomly chosen group of individuals, we make an assumption that the table mortality rates are suitable to this group and will remain constant, changing only with age, during the lifetimes of the members of the group. This assumption is not valid for the traditional life-settlement policies for a number of reasons, one of which is that future longevity improvements are not reflected in a traditional life table unless specific adjustments are made to the mortality rates. One of the outcomes is that the life expectancy e_x calculated based on the table may well be understated.

Below are some of the relationships between the mortality table functions that could be useful in the mortality analysis and mortality table construction and interpretation.

$$q_x = \frac{l_x - l_{x+1}}{l_x}$$

$$l_x = l_0 - \sum_{z=0}^{x-1} d_z$$

$$\frac{l_x}{l_0} = \prod_{z=0}^{x-1} p_z$$

In the stochastic application of a life table, the number of survivors to age x represented by l_x in the strictly deterministic case could be seen as having a binomial distribution with parameters l_0 and $S(x)$ if certain conditions are satisfied. One of the conditions is the independence of the mortality experience of individual members of the cohort. This assumption is strongly challenged for life settlement populations.

constructing a mortality table. Basic mortality table functions are also defined.

POPULATION MORTALITY TABLES

Census data is the initial point in constructing a general mortality table, which is also called a life table. An example of a life table is presented in Table 13.1, which shows an excerpt from the official US Life Table for the Total Population.

The table starts at the age of 0, the newborns, with 100,000 individuals. It then traces the number of deaths each year that reduce the surviving population. For example, the number of those dying between the ages of 0 and 1 is 680, reducing the surviving population at age 1 to 99,320. The mortality rate expressed as the probability of dying between the ages of 0 and 1 is then 0.68%. The table shows a more precise number for the mortality rate since the number of deaths between ages x and $x + 1$ is rounded, while the probability is not. The excerpt from the life table skips ages above 10 and gets immediately to the age-65-and-over category, which presents most interest to investors in mortality-linked securities and in particular in life settlements. Since the table is presented for illustrative purposes only, some of the other age ranges are also omitted. Out of the initial 100,000, according to the table only 9,419 survive to the age of 95; 1,873 of them do not survive to the age of 96, corresponding to the mortality rate 19.89%.

The life expectancy at the age of 95 is shown as 3.6 years. This is the

number of interest to investors in life settlement securities, but it cannot be used without adjustments, even if we were to assume that this particular table is applicable to a specific life settlement security. Elsewhere we also discuss why the widely used life expectancy parameter, when taken in isolation, is an inappropriate measure for analysing life settlement investments despite the fact that it is being used this way by some investors.

The table includes a column for the number of person-years lived between the ages x and $x + 1$. For ages between 95 and 96, this number is 8,482 out of the 9,419 surviving to age 95. There are significant differences in the way mortality tables are constructed, resulting from the purpose of a specific table, population segment for which the table is constructed, data used in determining mortality rates and methodology utilised in building the table from the available data. For this specific life table, the number of person-years lived between the ages of 95 and 96 (8,482) is the simple arithmetic average of those alive at the ages 95 and 96, implying that the calculation of this parameter is based on the assumption that a person dies in the middle of a year, in this case at 95.5 years of age.

Next, we further demonstrate how mortality tables could differ significantly, depending on which population they represent, their purpose and the way they are constructed.

The final comment about this specific table refers to ages 100 and over, for which mortality rates by year have not been calculated and only the aggregate number is shown, due to the low credibility of the data for older ages in general. The limited credibility of mortality data is a problem for all older ages, with ages 100 and over representing an extreme case. Mortality at ages over 65, which are of most interest for life settlements and similar securities, is much more uncertain than at younger ages. Mortality applicable to the subsets of the population of life settlements is even harder to determine: the statistical data sample is not big enough to be assigned a high level of credibility; projected changes in mortality in future years play a greater role in life settlements; and there are specific difficulties in determining the split by life settlement population subset.

Table 13.2 presents an excerpt from the official US Social Security Period Life Table that shows only older ages of most interest to investors in mortality and longevity risk. The table differentiates between males and females; and it can be seen that mortality differences by sex could be significant.

While based on substantially the same statistical data as Table 13.1, the table is constructed differently since it is intended to serve a different purpose. It is important to note that an attempt is made to calculate

Table 13.1 US Life Table for the Total Population (excerpt)

Age	Probability of dying between ages x to $x+1$	Number surviving to age x	Number dying between ages x to $x+1$	Person-years lived between ages x to $x+1$	Total number of person- years lived above age x	Expectation of life at age x
	q_x	l_x	d_x	L_x	T_x	e_x
0–1	0.006799	100,000	680	99,403	7,783,712	77.8
1–2	0.000483	99,320	48	99,296	7,684,309	77.4
2–3	0.000297	99,272	29	99,257	7,585,013	76.4
3–4	0.000224	99,243	22	99,232	7,485,755	75.4
4–5	0.000188	99,220	19	99,211	7,386,524	74.4
5–6	0.000171	99,202	17	99,193	7,287,313	73.5
6–7	0.000161	99,185	16	99,177	7,188,119	72.5
7–8	0.000151	99,169	15	99,161	7,088,943	71.5
8–9	0.000136	99,154	14	99,147	6,989,781	70.5
9–10	0.000119	99,140	12	99,134	6,890,634	69.5
65–66	0.014473	83,114	1203	82,513	1,553,230	18.7
66–67	0.015703	81,911	1286	81,268	1,470,718	18.0
67–68	0.017081	80,625	1377	79,936	1,389,450	17.2
68–69	0.018623	79,248	1476	78,510	1,309,513	16.5
69–70	0.020322	77,772	1580	76,982	1,231,004	15.8
85–86	0.085898	38,329	3292	36,683	261,765	6.8
86–87	0.093895	35,037	3290	33,392	225,082	6.4
87–88	0.102542	31,747	3255	30,119	191,690	6.0
88–89	0.111875	28,491	3187	26,898	161,571	5.7
89–90	0.121928	25,304	3085	23,761	134,673	5.3
95–96	0.198875	9,419	1873	8,482	33,889	3.6
96–97	0.214620	7,545	1619	6,736	25,407	3.4
97–98	0.231201	5,926	1370	5,241	18,671	3.2
98–99	0.248600	4,556	1133	3,990	13,430	2.9
99–100	0.266786	3,423	913	2,967	9,440	2.8
100+	1.000000	2,510	2510	6,473	6,473	2.6

Source: National Vital Statistics Reports 56(9), December 28, 2007

Excerpted from the “Life Table for Total Population: US, 2004”. Based on the final numbers of deaths for the year 2004, postcensal population estimates for the year 2004, and data from the Medicare programme of the Centers for Medicare and Medicaid Services. The tables are not cohort tables. The tables reflect general population mortality and differ from the mortality of the insured population. Life settlement mortality could differ even further.

mortality rates up to the age of 120, even though the data is so sparse that a number of assumptions have been made to determine the rates. None of the tables reflects the insured population that has mortality characteristics different from those of the general population.

Table 13.2 US Social Security Period Life Table (excerpt)

Exact age	Male			Female		
	Death probability	Number of lives	Life expectancy	Death probability	Number of lives	Life expectancy
41	0.002629	95,294	36.36	0.001581	97,418	40.52
42	0.002863	95,044	35.46	0.001732	97,264	39.58
43	0.003127	94,772	34.56	0.001891	97,096	38.65
44	0.003418	94,475	33.67	0.002059	96,912	37.72
45	0.003732	94,153	32.78	0.002244	96,713	36.80
46	0.004067	93,801	31.90	0.002441	96,496	35.88
47	0.004424	93,420	31.03			34.96
48	0.004805	93,006	30.17			34.06
49	0.005208	92,560	29.31			33.15
50	0.005657	92,077	28.46			32.25
51	0.006134	91,557	27.62			31.35
52	0.006595	90,995	26.79			30.46
53	0.007027	90,395	25.96			29.57
54	0.007457	89,760				28.68
55	0.007921	89,090				27.81
56	0.008467	88,385				26.94
57	0.009121	87,636				26.07
58	0.009912	86,837	21.92			25.22
59	0.010827	85,976	21.13			24.37
60	0.011858	85,045	20.36			23.53
61	0.012966					22.71
62	0.014123					21.89
63	0.015312					21.08
64	0.016567					20.29
65	0.017976	79,190	16.67			19.50
66	0.019564	77,766	15.96			18.72
67	0.021291	76,245	15.27			17.95
68	0.023162	74,621	14.59			17.19
69	0.025217	72,893	13.93			16.45
70	0.027533	71,055	13.27			15.72
71	0.030131					15.01
72	0.032978					14.31
73	0.036086					13.62
74	0.039506					12.95
75	0.043415					12.29
76	0.047789					11.64
77	0.052464	54,652	9.14			11.01
78	0.057413	51,785	8.62			10.40
79	0.062789	48,812	8.11			9.80
80	0.068836	45,747	7.62			9.22

mean estimate of the number of years in the remaining lifespan of a male at age 51


number of males reaching the age of 58 out of the 100,000 total at birth

probability of a male alive at age 65 dying before reaching the age of 66

Death probability in this table is *not* a straightforward outcome of the annual changes in the number of lives. The tables are different from the US Life Tables prepared by the Division of Vital Statistics of the US Department of Health and Human Services based on substantially the same data.

Table 13.2 Continued

Exact age	Male			Female		
	Death probability	Number of lives	Life expectancy	Death probability	Number of lives	Life expectancy
81	0.075724	42,598	7.15	0.053190	57,520	8.65
82	ages presumed to be most affected by longevity improvements		6.70	0.059279	54,460	8.11
83			6.26	0.066080	51,232	7.59
84			5.84	0.073685	47,847	7.09
85			5.45	0.082199	44,321	6.62
86	0.124164	26,116	5.08	0.091712	40,678	6.17
87	0.136917	22,874	4.73	0.102294	36,947	5.74
88	0.150754	19,742	4.40	0.113990	33,168	5.33
89	0.165704	16,766	4.09	0.126820	29,387	4.96
90	0.181789	13,988	3.80	0.140793	25,660	4.60
91	0.199019	11,445		0.155006	22,047	4.28
92	0.217396	9,167	data becomes very sparse and statistical credibility dangerously low		18,610	3.97
93	0.236906	7,174			15,406	3.70
94	0.257525	5,475			12,487	3.44
95	0.278031	4,065	2.68	0.226597	9,891	3.22
96	0.298111	2,935	2.52	0.245258	7,649	3.01
97	0.317432	2,060	2.38	0.263628	5,773	2.83
98	0.335655	1,406	2.25	0.281410	4,251	2.66
99	0.352438	934	2.13	0.298294	3,055	2.50
100	0.370060	605	2.02	0.316192	2,144	2.36
101	0.388563	381	1.91	0.335163	1,466	2.22
102	0.407991	233	1.81	0.355273	975	2.08
103	0.428390	138	1.71	0.376590	628	1.95
104	0.449810	79	1.61	0.399185	392	1.83
105	0.472300	43	1.52	0.423136	235	1.71
106	0.495915	23	1.43	0.448524	136	1.60
107	0.520711	12	1.35	0.475436	75	1.49
108	0.546747	6	1.26	0.503962	39	1.39
109	0.574084	3	1.19	0.534199	19	1.29
110	0.602788	1	1.11	0.566251	9	1.20
111	0.632928	0	1.04	0.600226	4	1.11
112	0.664574	0	0.97	0.636240	2	1.03
113	0.697803	data is so sparse that the same probability is assigned to both males and females		0.674414	1	0.95
114	0.732693			0.714879	0	0.87
115	0.769327			0.757772	0	0.80
116	0.807794			0.803238	0	0.73
117	0.848183	0	0.67	0.848183	0	0.67
118	0.890592	0	0.62	0.890592	0	0.62
119	0.935122	0	0.57	0.935122	0	0.57




The data represents the 2004 period tables published by the Office of the Chief Actuary of the US Social Security Administration. Data for ages 0–40 is not displayed. The life expectancy is calculated based on the assumption that the mortality rates of the period are experienced by survivors for the rest of the lives.

MORTALITY DYNAMICS

Figure 13.1 graphically illustrates mortality dynamics by showing how the number of lives out of the original 100,000 declines with age. The right-hand side of the chart is the one of primary interest in investing in mortality or longevity risk. It is also the area where changes are greater than at younger ages, and the sensitivity of results to assumptions is particularly pronounced.

Figure 13.2 displays the graph of the force of mortality μ_x showing its changes with age. The function steadily increases with the exception of some areas at younger ages shown on the insert. As expected, at older ages, which are of primary interest to investors, the force of mortality monotonously increases. The force of mortality could be used to model the variability of mortality and to perform stress testing of a portfolio of mortality-linked securities.


Some mortality models quantify deaths attributable to various causes, examples of which could be accidents, different types of cancer or cardiovascular diseases. The approach could be particularly useful in the analysis of life settlement investments. In these multiple decrement models, a simplifying assumption used in the initial analysis could be that the force of

Figure 13.2 Changes in force of mortality by age

mortality is constant over each year of age for each decrement. Analytical mortality models that have been used by actuaries include the Gompertz, Makeham, de Moivre, Weibull and other formulas for the force of mortality, each of which results in a different form of the survival function.

The modern approach, however, is to use stochastic modelling, which in most cases eliminates the need to postulate a specific analytical form for the force of mortality, or even to use the force of mortality in formulating a model. Examples of stochastic modelling of investment portfolios of mortality and longevity risk are provided in subsequent chapters. It is important to point out that stochastic models, while widely recognised as the best approach to mortality modelling, are not widely used by investors. The models could be very sensitive to assumptions and always benefit from a reasonability check. Such a reasonability check could be provided by using a simpler model based on one of the analytical formulas for the force of mortality.

Figure 13.3 is the graph of the $l_x \mu_x$, which under certain assumptions could be interpreted as representing the density of deaths. It is sometimes referred to as the curve of deaths. Again, investors are interested in the right-



hand side of the graph, but at older ages this function is not monotonous and has a relatively sharp maximum. For most populations in the US and Europe, including the standard risk class in life insurance underwriting, this maximum occurs between the ages of 75 and 85, at the inflexion point of the l_x function shown in Figure 13.1.

The life tables described above are called aggregate life tables because they are limited to a single mortality rate q_x for each age x . The insurance industry and investors in mortality products generally use mortality tables designed specifically for the segments of the population being insured or referenced. Segments of the population defined, for example, based on the perceived mortality risk, have their own distinct mortality levels and patterns, which are reflected in these mortality tables.

SELECT AND ULTIMATE TABLES

An individual applying for an insurance policy usually goes through the process of underwriting, most of which is focused on determining their medical condition in order to assign the applicant to the right risk class. The underwriting happens only once, before a policy is issued, and changes in the medical condition after the policy issuance are not considered.

Investors in life insurance policies also have the ability to underwrite a

policyholder using their own expertise in the field, or going the more common route of utilising the services of a life expectancy provider. Assessing the value of the policy is based primarily on the results of this underwriting. Investors can perform their underwriting at the point when the insurance policies have already been issued and the rates set by the insurance companies. This puts investors in a position of being able to differentiate among several risk levels within the same category of policyholders that was set by the insurance company at the point of its underwriting.

Starting with underwriting performed by the insurance industry, we can say that its purpose is to stratify the applicants based upon their expected mortality. Two individuals of the same age might or might not be judged to have the same mortality characteristics; consequently, they might or might not be assigned the same mortality table. An individual of age x , after having undergone the process of life insurance underwriting, will be judged to have a select mortality rate $q_{l[x]}$ different from the ultimate mortality rate q_x of the general population. The q_x is usually not the mortality rate of the aged- x population as a whole but rather the mortality rate of the subset of the general population with the same sex or other characteristic as the policyholder, but without taking into account other underwriting factors.

In the case of underwriting risk classes with better-than-average expected mortality $q_{l[x]} < q_x$, underwriting skill is based on the ability to assign the right mortality level on the basis of the information obtained in the underwriting process. Two different underwriting processes could result in two groups of insureds being initially assigned the same mortality rate q_x . While this mortality rate for both groups might be correct, it does not imply that both underwriting processes produce the same results. One could be judged superior to the other if it produces mortality rates distinct from the ultimate even many years after the individuals have been underwritten.

We label $q_{l[x]+t}$ the mortality rate of an individual aged $x+t$ who was underwritten t years ago at age x . If at the time of underwriting we had $q_{l[x]} < q_x$, we can also expect that $q_{l[x]+t} < q_{x+t}$. The underwriting effect wears off over time, and select mortality rates slowly revert to ultimate rates. After s years, $q_{l[x]+s} \cong q_{x+s}$. The number s , the smallest of such numbers for which this is true, is called the select period. In other words, for the duration of s years, select mortality is different from ultimate, after which the effects of underwriting wear off and the select mortality rate converges to the ultimate mortality rate. The ratio of the mortality rates, $q_{l[x]+s} / q_{x+s}$, as opposed to the difference between the two, is the right measure to use in determining when the selection has worn off. In practice, the selection period is chosen

so that it would be the same for all ages in a mortality table, or at least for a relatively wide range of ages. During this selection period, mortality rates are different from the ultimate rates for the same ages. In other words, for each age we have not a single mortality rate but rather a set of mortality rates for the number of years equal to the selection period, after which the rates revert to ultimate.

It means that the mortality rate, as well as each of the other mortality table functions, now becomes a two-dimensional array, where a row corresponding to the age $[x]$ has $s + 1$ elements including the ultimate mortality rate at the age $x + s$. Mortality table functions and survival functions in general each become functions of two variables, $[x]$ and t . In the case of whole ages, this bivariate function translates into a two-dimensional array represented by the two-dimensional select-and-ultimate mortality table. The last column of a select-and-ultimate table represents the ultimate mortality rates; these no longer have any effect of initial underwriting and correspond to the aggregate mortality table.

To illustrate the concept of a select-and-ultimate mortality table, Table 13.3 shows an excerpt from the 2008 Valuation Basic Table, the latest mortality table compiled by the Society of Actuaries for US insured lives. This table is described in greater detail later in this chapter. The selection period in this table is 25 years. For selection age 59, mortality rate is 0.2%. General, or ultimate, mortality rate at the same age is 0.6%. For the second year of the selection period, mortality rate is 0.3%. Mortality rate increases every year and reaches 6.8% for the last year of the selection period. After that, mortality rate is assumed to revert to the ultimate, which is equal to 7.7% for age 84, the first year after the selection period ends. For following ages, mortality rate is taken from the last column of the table.


CREDIBILITY THEORY APPROACH

Credibility theory approach could be useful in constructing a mortality table, especially when deciding to what degree the difference between actual mortality and that based on a mortality table warrants making adjustments to the table. The actual-to-expected (A/E) ratio is the parameter most commonly used in this type of analysis; it is notable that, in practical applications involving small data samples, choosing to analyse the A/E for mortality rates could produce results different from those based on choosing the analysis of A/E for life expectancies or the number of lives surviving at specific ages. This is a common situation in analysing life settlement populations, where even the life expectancy providers do not have sufficient data

Table 13.3 2008 Valuation Basic Table (excerpt from the Male SUN ANB)

Age [x]	Mortality rate (%)																									Age [x]+25
	Year following issue																									
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	
55	0.1	0.2	0.2	0.3	0.3	0.4	0.4	0.5	0.6	0.7	0.8	0.9	1.0	1.1	1.2	1.4	1.6	1.8	2.1	2.4	2.7	3.1	3.5	4.0	4.6	5.2
56	0.1	0.2	0.3	0.3	0.4	0.4	0.5	0.6	0.7	0.7	0.8	1.0	1.1	1.2	1.4	1.6	1.8	2.0	2.3	2.6	3.0	3.4	3.9	4.4	5.0	5.7
57	0.1	0.2	0.3	0.3	0.4	0.5	0.5	0.6	0.7	0.8	0.9	1.1	1.2	1.4	1.5	1.8	2.0	2.3	2.6	2.9	3.4	3.8	4.3	4.9	5.6	6.3
58	0.2	0.2	0.3	0.4	0.4	0.5	0.6	0.7	0.8	0.9	1.0	1.2	1.3	1.5	1.7	2.0	2.2	2.5	2.9	3.3	3.7	4.2	4.8	5.4	6.1	7.0
59	0.2	0.3	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0	1.2	1.3	1.5	1.7	1.9	2.2	2.5	2.8	3.2	3.6	4.1	4.7	5.3	6.0	6.8	7.7
60	0.2	0.3	0.4	0.5	0.5	0.6	0.7	0.9	1.0	1.1	1.3	1.5	1.7	1.9	2.2	2.5	2.8	3.2	3.6	4.0	4.6	5.2	5.8	6.6	7.5	8.6
61	0.2	0.3	0.4	0.5	0.6	0.7	0.8	1.0	1.1	1.3	1.4	1.6	1.9	2.1	2.4	2.7	3.1	3.5	4.0	4.5	5.1	5.7	6.4	7.3	8.3	9.5
62	0.2	0.3	0.4	0.6	0.7	0.8	0.9	1.1	1.2	1.4	1.6	1.8	2.1	2.4	2.7	3.1	3.5	3.9	4.4	5.0	5.6	6.3	7.1	8.2	9.3	10.6
63	0.2	0.4	0.5	0.6	0.7	0.9	1.0	1.2	1.4	1.6	1.8	2.0	2.3	2.6	3.0	3.4	3.8	4.3	4.9	5.5	6.2	7.0	8.0	9.2	10.5	11.9
64	0.3	0.4	0.5	0.7	0.8	1.0	1.1	1.3	1.5	1.8	2.0	2.3	2.6	3.0	3.3	3.8	4.3	4.8	5.4	6.1	6.8	7.9	9.0	10.3	11.8	13.3
65	0.3	0.4	0.6	0.8	0.9	1.1	1.3	1.5	1.7	2.0	2.2	2.6	2.9	3.3	3.7	4.2	4.7	5.4	6.0	6.7	7.8	8.9	10.2	11.6	13.2	14.8
66	0.3	0.5	0.7	0.8	1.0	1.2	1.4	1.6	1.9	2.2	2.5	2.8	3.2	3.7	4.1	4.7	5.3	5.9	6.6	7.7	8.8	10.1	11.6	13.1	14.8	16.5
67	0.4	0.5	0.7	0.9	1.1	1.3	1.6	1.8	2.1	2.4	2.8	3.2	3.6	4.1	4.6	5.2	5.9	6.6	7.6	8.8	10.0	11.5	13.1	14.8	16.5	18.3
68	0.4	0.6	0.8	1.0	1.3	1.5	1.8	2.0	2.4	2.7	3.1	3.5	4.0	4.6	5.1	5.8	6.5	7.5	8.7	10.0	11.4	13.1	14.8	16.5	18.3	20.1
69	0.4	0.7	0.9	1.1	1.4	1.7	2.0	2.3	2.6	3.0	3.5	3.9	4.5	5.1	5.7	6.4	7.5	8.6	9.9	11.4	13.0	14.8	16.5	18.3	20.2	22.1
70	0.5	0.8	1.0	1.3	1.6	1.9	2.2	2.5	2.9	3.4	3.8	4.4	5.0	5.6	6.3	7.4	8.6	9.9	11.4	13.0	14.8	16.5	18.3	20.2	22.1	23.7

Source: Society of Actuaries

Figure 13.4 Life expectancy at 65 years of age by sex, US

Based on the data in *Health, United States, 2008* published by the National Centre for Health Statistics, Centres for Disease Control and Prevention, US Department of Health and Human Services. Death rates used to calculate life expectancies for 1997–1999 are based on postcensal 1990-based population estimates; life expectancies for 2000 and beyond are calculated with death rates based on census 2000. Deaths to non-residents were excluded beginning in 1970. Original data sources include National Vital Statistics Reports, US Life Tables.

samples observed for sufficient time periods to assign significant credibility to the results. The small sample size makes it difficult to determine the degree of accuracy of the life settlement mortality tables assembled by some of the life expectancy providers.

LONGEVITY IMPROVEMENTS

Most of the countries in the developed world are seeing people live longer than in the past. This trend has continued for many decades, if not centuries. Focusing on the past 20 years, we can say that longevity improvements have been unexpectedly high. Longevity improvements are attributed to better living conditions and improved medical care. Continuing medical advances are expected to contribute to longevity improvements in the foreseeable future. The degree of the improvements is a topic of heated discussions; there is no consensus at this point.

Figure 13.4 shows the change in life expectancy at the age of 65 in the US over the past four decades. Life expectancy has steadily climbed for both males and females, with slightly greater longevity improvements for males than females. This is a common pattern observed in most countries in the developed world as well as in the risk subsets of the insured populations.

Longevity improvements of the population as a whole do not necessarily translate into longer lifespans for everybody. The epidemic of obesity in the US has had the effect of lowering expected longevity for many. Most projections anticipate that it will continue to have a negative impact, possibly to a greater degree, but this impact will be on average more than offset by the other factors affecting longevity in the positive way.


Figure 13.5 presents life expectancy for males and females in the UK based on actual mortality rates and projected mortality rates from the 2006-based population projections by the Government Actuary's Department. Again, we can see steady improvements in mortality rates over the years for both males and females. These life expectancies, however, are different in that they have been calculated making specific assumptions as to the future mortality rates, as opposed to assuming that current age-adjusted mortality rates will remain constant in the future and can be used in the calculation of life expectancy. Mortality rates are expected to decrease in the future even more.

The Government Actuary's Department of the UK calculates these expected improvements and builds the base scenario, termed the principal projection. Around this base case, two additional projections are built: the high life expectancy variant projection and the low life expectancy variant projection. All three assume longevity improvements; the low life expectancy variant projection has longevity improvements for several years, after which life expectancy stays at the same level.

Actuaries and demographers have been producing estimates of longevity improvements for many years. The need to develop a view on future changes in the mortality rates is important for most life insurance and annuity products as well as for pension funds and social security systems. By and large, the projections have so far turned out to underestimate the magnitude of longevity improvements. The actual mortality rates have been decreasing at a fast pace; to what degree this trend will continue remains an open question.

Longevity improvements have not had the same impact on all segments of the population. Figure 13.4 shows the difference in longevity improvements between males and females in the general population. There is an even greater difference between the insured population and the general population. Within the insured population, the rate of change in longevity improvements has differed sharply by risk category. Knowing these differences is critical in the analysis of life settlements and similar securities, in particular when the life expectancy is relatively long. There is also credible evidence that not only do mortality rates themselves vary with policy face

Figure 13.5 Cohort expectation of life at age 65, UK



value, but the mortality improvements for the insured population differ by policy face value as well.

While definitive conclusions as to the degree of future longevity improvements by category cannot be achieved, the significant body of collected statistical data and external inputs do permit the development of reasonable models of future mortality. The external inputs include medical data and

expectations of improvements in the treatment of major diseases. The models have to be stochastic due to the inherent uncertainty involved in such projections.

It is also possible to determine sets of scenarios of future mortality rates and assign probabilities to these scenarios. In the context of portfolio management, the use of scenario testing is not only possible but required, in particular in the context of risk management of portfolios of mortality risks. Unfortunately, very little of the body of data that has been developed is being used in determining life expectancy and mortality levels for investment portfolios of life settlements. Poor modelling of portfolios of life settlements also prevents effective use of mortality hedging instruments in managing portfolio mortality risk.

LEE-CARTER AND RELATED METHODS

The Lee–Carter method is an example of mortality models that are particularly useful in analysing older-age longevity. In the simplified Lee–Carter model, the force of mortality is determined the following way

$$\ln \mu_{xt} = \alpha_x + \beta_x k_t + \varepsilon_{xt}$$

where μ_{xt} is the force of mortality at age x in year t , ε_{xt} is the parameter introducing randomness, and all the other parameters are estimated from the available data. α_x describes age-specific mortality at age x . k_t describes the general level of mortality in year t . β_x reflects the sensitivity of the mortality at age x to the changes in general mortality level in year t .

Parameter β_x is assumed to be constant over time. This assumption is important in that it is a source of potential error and also a way to simplify model parameterisation. A common assumption is that $\sum k_t = 0$ and $\sum k_t^2 = 1$, resolving the problem of non-uniqueness of the model parameterisation.

If we assume that ε_{xt} is normally distributed, parameters k_t and β_x can be estimated using the maximum likelihood approach or another method. Lee and Carter also suggest the use of singular value decomposition to determine the optimal values as well as performing the second stage estimation, or re-estimation, once the optimal values have been arrived at in the first stage.

Parameter k_t is typically modelled as a random walk with a constant drift; that is,

$$k_{t+1} = k_t + d + e_{t+1}$$

where d is the drift parameter and e_t is the random term. e_t is assumed to be normally distributed with the mean of zero, $e_t \sim N(0, \sigma^2)$.

The Lee–Carter method is based on extrapolating existing data and thus might perform poorly as the time horizon increases. The method is further described in the chapters dealing with specific applications of longevity modelling in insurance-linked securities.

MARKOV PROCESS OF MORTALITY AND MORBIDITY

While the use of mortality tables in a stochastic environment could by itself be seen as a Markov process application, a more advanced Markov process-based model of mortality and ageing could be extremely useful in analysing mortality, in particular that of older ages.

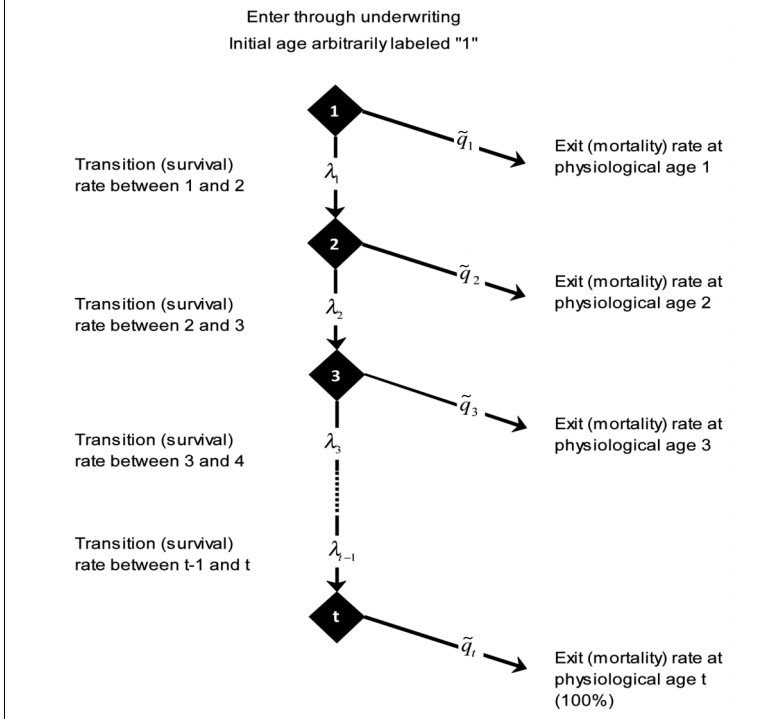
Even in its continuous form, mortality could be modelled as a Markov process. Brownian motion with a drift, not necessary constant, could also be incorporated in this framework. Many of the models explicitly using Markov process to model mortality produce results at older ages that are not always easily explainable. Mortality plateaus at older ages are typical in such models. There is no consensus on how to introduce a mortality correction dealing with this potential problem and whether it is necessary at all.

Many theoretically sound models run into problems in practical applications such as analysing insurance-linked securities. The problems often have to do with parameter risk of the models, which in turn is usually the result of incomplete information.

Mortality modelling using physiological age

The concept of physiological as opposed to calendar age can be used to describe the combination of the actual age and health condition of an individual. The two factors determine both the probability of dying and the speed of progression to the next physiological age. There could be a significant difference between the actual age of a person and their physiological age. Physiological age is essentially the calendar age adjusted for the relative health condition of the individual; using physiological age in mortality modelling introduces a number of parameters into the model. In a simple but effective framework¹ illustrated in Figure 13.6, mortality and ageing are modelled by a finite-state continuous-time Markov process of moving from one physiological state to another, with an exit state associated with mortality.

Transition and exit processes are described by the parameters λ_k and \tilde{q}_k . λ_k measures the ageing between the physiological ages of k and $k+1$ and could be seen as a rate or force of transition from the physiological age k to

Figure 13.6 Progression in physiological age

the age $k+1$. \tilde{q}_k measures the probability of an individual at the physiological age k dying before reaching the age of $k+1$.

An actuary would be strongly tempted to interpret \tilde{q}_k as q_k , the traditional mortality rate at age k , and λ_k as its complement, the survival rate between the ages of k and $k+1$, that is, the rate of p_k . This interpretation would be correct only if we were also interpreting ages k and $k+1$ as the actual as opposed to the physiological ages of the individual, which would then be seen as a traditional mortality table transition.

In the framework described here, the ages are physiological and not calendar, bringing about a very different interpretation of the process. \tilde{q}_k is separate and distinct from the mortality rate q_k as it is traditionally defined. (For simplicity, the physiological ages in Figure 16.6 start at age 1 at the point of underwriting, as opposed to the real physiological or calendar age.)

Physiological ages are not equidistant as calendar years are. In progression from one physiological age to the next, the ageing process follows its

own rules based on a number of parameters, the most important of which is the relative health of the individual.

The transition matrix of the ageing process with one exit state could be expressed in a simplified form as

$$\hat{T} = \begin{pmatrix} -\lambda_1 - \tilde{q}_1 & \lambda_1 & 0 & \dots & 0 \\ 0 & -\lambda_2 - \tilde{q}_2 & \lambda_2 & \dots & 0 \\ 0 & 0 & -\lambda_3 - \tilde{q}_3 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & -\tilde{q}_t \end{pmatrix}$$

The challenge is in determining the set of parameters λ_k and \tilde{q}_k that would describe the process with a sufficient degree of accuracy. This determination is closely tied to the question of the definition of physiological ages. As there are many ways to define physiological ages within the same overall framework, there are also many ways to specify the parameters of the transition matrix.

One way of defining physiological age is to keep all λ_k 's constant; that is, to have $\lambda_k = \lambda$ for all k . This is logical and creates a very clear picture of physiological ages. At the same time, it makes it more difficult to define \tilde{q}_k 's in a way that incorporates existing mortality statistics based on traditional calendar ages. The difference between physiological and calendar ages is particularly pronounced at older ages, which happen to be of most interest in modelling mortality in life insurance-linked securities.

Moving from population to individual mortality

While the standard approach in using physiological age focuses on population mortality and relies on averaged population statistics, there is no reason why the framework cannot be applied to subsets of the overall population meeting certain underwriting criteria.

Parameters \tilde{q}_k can be decomposed into components reflecting various causes of mortality, including accidents, specific types of cancer, cardiovascular disease, HIV / AIDS and many others. Then we have $\tilde{q}_k = a_1 g_1^k + a_2 g_2^k + \dots + a_m g_m^k$, where g_i^k is the value of the i th factor affecting mortality at the physiological age of k (where k is, again for simplicity, defined so that the physiological age at underwriting or other entry point is 1. Each of the m coefficients a_i measures the contribution of the corresponding factor to the overall mortality. Recognising that factors g_i^k play a different role at different physiological ages, weights they are assigned could also vary by physiological age:

$$\tilde{q}_k = \sum_{i=1}^m a_i^k g_i^k$$

Expressing the \tilde{q}_k 's mortality parameters this way implicitly assumes that the contributions of factors g_i^k are additive. The assumption of additivity is common in traditional life insurance underwriting, which is based on actual calendar ages. The assumption simplifies the analysis but is a source of error, in particular at older ages. When physiological ages are used, it could be beneficial to use factors that are multiplicative instead of additive. Ideally, a combination of the multiplicative and additive approaches would be utilised to arrive at the mortality parameters \tilde{q}_k 's.

The assumption that λ_k is the same at all physiological ages provides a simple way to define physiological ages, but this framework does not always represent the best way to model mortality. Having λ_k 's vary by age could improve the model and allow for a more accurate determination of the \tilde{q}_k 's parameters. λ_k could also be modelled as a random variable.

Age transform

Modelling done in terms of physiological ages necessitates transforming physiological ages back into calendar ages. The definition of physiological age is key to constructing an appropriate transform function. The assumption of all λ_k 's being the same at all physiological ages or at certain ranges of physiological ages could simplify modelling but does not make it easier to translate physiological ages back into calendar ages.

The level of precision in the age translation is a consideration in deciding whether the physiological age-based model should be used instead of the more traditional model based on actual calendar ages. The precision and accuracy that a physiological age-based model may provide could all be lost if the age transform function introduces significant error.

DIRECT AGE TRANSFORM IN MORTALITY MODELLING

There is a way to use the concept of the age transform function in the more traditional analysis based on actual instead of physiological ages. Physiological ages are used too, but in an indirect fashion.

An underwriter can determine that an insured, due to their medical condition, lifestyle and other factors, has a mortality propensity equal to that of an individual two years younger but having a different health profile. The comparison calls for using mortality functions corresponding to that other hypothetical individual. In simple terms, this determines the choice of a mortality table to be used and the "entry point" when using this table.

Mortality underwriting of this kind is different from the model where debits and credits are assigned but the choice of the table, and in particular the entry point, is fixed.

The important point to make is that, in this approach, not only is the current mortality rate matched to that of a different “entry point” and possibly a different table, but the longer-term mortality functions should match as well. In other words, the expected mortality rate in 10 years should also be the same. If such correspondence cannot be established, the approach fails and needs to be modified.

Underwriting should be based on a multifactor model that takes into account many of the factors affecting mortality, similar to the way it would be done in modelling physiological ages. The underwriting could be based on the structural and functional variables that measure physiological ageing in the ageing kinetics framework or, when sufficient information is not available, the more traditional underwriting variables.

As with the select period used in construction of mortality tables, there is a tapering off of the underwriting effect with time. This leads to a reversion to the ultimate mortality rates. Depending on the way underwriting is performed and the choice of mortality tables and modifiers, these ultimate rates are likely to be neither the ones from the table chosen nor the ones from the table corresponding to the actual chronological age of the individual.

This type of modelling is most valuable in underwriting for older ages, which are prevalent in life settlements and some other insurance-linked securities.

MORTALITY AND LONGEVITY SHOCKS

Most models of mortality and longevity that include randomness and resulting variability do not allow the modelling of shock events. In the context of insurance-linked securities, the two most important types of shock events are these.

- **CATASTROPHIC MORTALITY EVENTS:** An example of such an event would be a pandemic resulting in a sudden jump in mortality rates. It could be caused, for example, by a swine or bird flu spreading among human populations and causing a high level of fatalities. This would have a jump effect over a relatively short period of time followed by a reduction of mortality rates to the levels comparable to those before the catastrophic event. Extreme mortality bonds are an example of an instrument transferring the risk of catastrophic mortality events from the insurance industry to investors.

- **PERMANENT LONGEVITY SHIFT:** An example of such a shift would be the unexpected magnitude of sustainable longevity improvements. It could have the form of a true shock such as longevity improvements due to a significant medical advance in the treatment of cancer. It could also be a “slow shock” resulting from a combination of forces leading to the same effect of permanent longevity improvements. Unexpected longevity shifts are important to many insurance-linked securities and need to be carefully analysed by investors.

Modelling mortality and longevity shocks is best accomplished in the simulation framework by incorporating a jump component into the mortality rates. This approach works for both temporary and permanent jumps but is highly dependent on the model assumptions.

Mortality and longevity shocks are described in more detail in other chapters.

CONCLUSION

This chapter introduced some of the key concepts in the modelling of mortality and longevity risk. Other chapters build on these concepts and describe some additional tools for mortality and longevity analysis.

There exist numerous models of mortality and longevity, with no single model that could be used in all circumstances. The choice of model depends, to a significant degree, on a specific type of product being analysed and the type of mortality risk it involves.

The second and sometimes overriding constraint on the choice of mortality risk model is data availability. While it is possible to develop a model incorporating numerous factors affecting mortality, limited data serves to reduce the credibility of any results produced by such a model. Proper parameterisation of such a model is usually impossible. For any type of model, we have to find the right balance between the desire to incorporate a greater number of model parameters and the credibility of estimates of these parameters based on the available data. It is important to choose the model that will properly capture the relevant mortality risk and utilise the data in a way that produces credible results.

Incorporating trends has been a challenge in modelling all types of insurance risk. Mortality has seen an important trend in the form of longevity improvements over a significant period of time. As with any trend, increase in longevity introduces additional uncertainty since the degree to which the observed trend could be extrapolated into the future is never known. The

long projection periods involved in the analysis of many types of life-insurance-linked securities contribute to this uncertainty.

Mortality shocks are important in valuing certain types of insurance-linked securities. Traditional mortality models do not incorporate mortality shocks, nor are shocks usually present in the historical data used in model construction and parameterisation. Since some insurance-linked securities are designed to transfer the risk of mortality shocks to investors, this risk has to be properly modelled to value these securities and to manage investment portfolios that contain them.

The use of a stochastic approach is essential to modelling mortality. Mortality models have historically developed in a deterministic context, which does not allow them to fully capture the inherent variability of the key mortality and longevity characteristics. Many insurance-linked securities are concerned primarily with this volatility and can be properly analysed only within a stochastic framework. Even when the primary mortality risk being transferred is not that of potential changes in mortality, there is a need for a stochastic framework, since investors have smaller portfolios of the risk than insurance companies and thus are likely to experience greater variability of investment results.

The field of mortality and longevity modelling is broad and continues to expand, and the information in this chapter provides but a foundation for exploring it further. It is noteworthy, however, that many of the models currently used by practitioners in pricing mortality-linked securities are quite simplistic. It creates the need for greater sophistication in modelling mortality and longevity, which will lead to advances in modelling life insurance-linked securities.

¹ The framework borrows the key elements of the approach proposed by X. S. Lin and X. Liu (2007) but provides a slightly different interpretation of the parameters describing the ageing process in the context of mortality modelling.

Valuation of Life Settlements and Other Mortality-Linked Securities

MODELLING INVESTMENT PERFORMANCE OF LIFE SETTLEMENTS

The main focus of this chapter is investor analysis of life settlement securities. Concepts described here could also be applied in investment analysis of other types of insurance-linked securities. The chapter builds on the ideas and theories described earlier and provides the foundation for further exploration of these ideas.

Analysing investments in life insurance settlements presents unique challenges to an investor not well familiar with life insurance. The nature of the risks and uncertainty differs significantly from those involved in almost any other type of investment, and none of the standard fixed income or other analytical approaches are applicable. On the other hand, insurance professionals seem quite comfortable in identifying the risks and building pricing models for life settlement mortality.

Both the standard investment approach and that utilised by some life insurance professionals work, but work poorly. Analysis of life settlements requires both a thorough understanding of insurance and a capital-markets perspective. Only a combination of the two could allow us to perform sound analysis of life settlement investments. The general framework for such analysis has been developed, but there is a clear need to improve on the current methodologies. More importantly, part of the life settlement marketplace is still using simplistic approaches to analysing these securities, and we have to realise that this could lead to mispricing.

An individual life settlement for a policyholder aged x could be analysed by projecting cashflows and using the standard net present value (NPV) approach

$$NPV_x = \sum_{t=0}^{\text{lim age} - x} \frac{{}_t p_x q_{x+t} \left[\frac{\text{Policy Benefit} - \text{Claim Expense}_t}{(1 + IRR)^{t+1}} \right]}$$

$$- \sum_{t=0}^{\lim age - x} \frac{{}_t p_x \left[\frac{Premium_t - Maintenance}{Expenses_t} \right]}{(1 + IRR)^t} - Price - \frac{Underwriting}{Expenses}$$

The formula for NPV is simplified; in reality, deaths could occur at any point during the period; the death benefit could be reduced by outstanding loans; and other adjustments might need to be made. It is assumed that brokerage and related fees are included in either price or underwriting expenses. Underwriting expenses are also intended to include the expense of setting up a trust if required. The premium payment pattern could differ. Maintenance expenses include monitoring expenses and administration costs. Claim expense is the cost of obtaining death benefits from the insurance company, as well as any associated expenses such as those related to any trust arrangement that might need to be terminated. The expense level could change over time for a variety of reasons, the simplest being inflation.

A critical choice is that of the mortality table applicable to the specific policyholder. The values of ${}_t p_x$ and q_x are specific to the insured. The result is strong sensitivity to the choice of life expectancy and mortality rates. A seemingly small change in mortality rates could produce a sizable change in the net present value of a policy and the price an investor would be willing to pay for it. Additional simplifications include the assumption that the internal rate of return (IRR) demanded by the investor does not differ by time period. In reality, a higher IRR would usually be required for longer-term investments.

Finally, the NPV formula above is the expression for the expected value and does not take into account variability of results around the mean. Of interest is the variability of the NPV, including the probability of its being in the negative territory, or, in a different approach, the variability of the rate of return that could be realised on this investment. Greater uncertainty and variability lead to a higher rate of return demanded by investors. The stochastic approach is described in the following chapters in the context of portfolio modelling.

A more comprehensive framework also takes into account the cost of having funds available in the future, associated with the need to ensure that premium and expense payments are made as long as the insured is alive (which itself is a random variable). We also need to consider the reinvestment risk in cases where the death payment was made earlier than expected. A number of additional factors play a role in portfolio management, as

many policies are analysed as part of a total portfolio. Depending on the composition of the portfolio, a given policy might or might not be a welcome addition.

LIFE EXPECTANCY

Prediction is very difficult, especially if it's about the future.

Niels Bohr

Life expectancy – or LE, as it has become known in the life settlement industry – is defined as the mean future lifetime of a person until death. It is important to emphasise that LE is the expected (average) value as opposed to the median or the mode. The term LE has been incorrectly interpreted as the median value quite a number of times. The use of the word “expected” in the definition of LE has also led some unsophisticated investors to view LE almost as a deterministic measure, where an “LE of 10 years” is taken to mean that the person will die in 10 years or very close to 10 years. This perception has led to underestimating the inherent variability of an individual's lifespan. Many cashflow projections presented to potential investors to demonstrate expected investment return have shown only the “base” scenario, where insureds die in exactly the number of years assigned to them by the LEs.

The need to value insurance policies in the context of life settlements has led to the appearance of a cottage industry of so-called LE providers. LE providers are firms that, based on the medical history and other information about insured individuals, estimate their LE. In the beginning, the only output of their analysis was a single number for LE such as 84 months. As the level of sophistication has grown, most LE providers now also make available a probabilistic view of mortality. Such a view could include a mortality table that an LE provider believes to be applicable to a particular individual.

LE determination

How are LEs determined? The answer to this question depends on the specific LE provider. The concept of LE as a key measure of a policy's value originated in viaticals, where LE was usually very short. LE was then used as a single measure of mortality; no mortality tables were used because the estimates were based entirely on medical records and often a highly subjective analysis of short-term longevity. While the level of sophistication has grown significantly since those early days of the market, some LE providers still use a somewhat backwards approach for life settlements with relatively

short LEs: they start with the single-number LE and then attempt to find corresponding mortality rates, instead of more properly determining the mortality rates and using them to calculate the LE.

Most LE providers employ medical underwriters with experience working in the underwriting departments of life insurance companies. Often they have some medical background that could be supplemented by the expertise provided by medical practitioners on staff or acting in a consulting capacity. In the ideal case they have strong expertise in old-age mortality and are well familiar with major diseases afflicting the 65-and-older population, as well as with the longevity of younger people who have serious medical conditions. The analysis could lead, for example, to placing the insured in a certain category for which the LE provider has developed a mortality table or LE point estimate. This mortality table or point estimate could be further adjusted using additional information specific to the insured.

Medical underwriting adds most value for insureds with shorter LEs, typically those suffering from a serious illness. For those whose health is less impaired, LEs are longer and the approach of using modifications to one of the standard mortality tables plays a greater role. In general, medical underwriting tends to drive the determination of mortality rates in the first years, with these rates reverting more to standard mortality tables later on. We could interpret this approach as using a weighted average of the underwriting and mortality-table approaches, with the weights assigned to medical underwriting decreasing every year. Given that some LE providers started developing their methodologies for purposes of viatical settlements, there is occasionally too much emphasis, even for longer LEs, on determining LE point estimates directly from the evaluation of medical conditions and then backing into a mortality table to determine mortality rates. This approach could lead to significant errors.

There is intense competition among LE providers. Each uses its own methodology; however, the differences in methodology are difficult to analyse and many details are not disclosed. It is also important to note that, in determining LE, a significant degree of judgement is often used, as opposed to following a purely formulaic approach. Comparison among LEs determined by different LE providers is difficult. Most investors require LE estimates from more than one source before making an offer for a policy. Some could use average LEs obtained from three sources; or could discard the outlier and average the remaining two. Others might pay particular attention to the outlier and try to determine the reasons why one estimate

differs significantly from the others. This determination requires a certain degree of underwriting expertise that few investors possess. The conservative ones would go with the longer LE. Still others have formed opinions on the quality of the work performed by individual LE providers and would treat estimates from some LE providers as more credible than the results produced by others. One way to do it would be by calculating a weighted average of LEs, with weights corresponding to the credibility of specific LE providers. We leave aside the question of whether averaging LEs or averaging corresponding mortality rates is the correct approach.

METHODOLOGY CHANGES IN THE CALCULATION OF LIFE EXPECTANCY

One of the things that I think we have learned is that we should all be very careful about making predictions about the future.

Bill Clinton

Comparing one LE provider to another is difficult in the absence of full disclosure of the methodology employed and the type of judgement calls made. It is particularly difficult when the methodologies have been changing.

Some attempts at comparing actual to expected mortality rates have been made. Results seem to indicate that certain LE providers have consistently underestimated LEs. One LE provider known for conservatism and producing longer life expectancies and lower mortality rates has had an actual-to-expected (A/E) ratio surprisingly close to 100%. This finding supports the view that the “conservatism” of this LE provider is in fact realism. Unfortunately, the data is limited, leading to the credibility of actual-to-expected comparisons being limited as well. In addition, the time horizon over which the comparison can be performed is too short to allow us arrive at truly meaningful conclusions.

In 2008, some LE providers made significant changes to their overall methodologies. The effect was to bring the results produced by different LE providers closer to each other, which raises a question of a prior bias in the determination of life expectancies. This change left some investors reeling, as all their profitability projections were based on what are now called “old LEs”. Policies that were purchased based on these old LEs remain in many investment portfolios and will stay there for many years to come. It is doubtful that many of these policies have been repriced for portfolio valuation purposes, leaving the possibility that a number of investment portfolios might have overstated net asset values.

The introduction of the 2008 Basic Valuation Table of mortality described

below made most LE providers reassess their methodologies, even in cases where no systematic bias was found. It will likely become the main reference point for LE providers in estimating LEs, even where proprietary methodologies based on specific adjustments to other mortality tables have been developed.

UNDERWRITING CONCEPTS

Here we review some of the life insurance underwriting concepts that have applications to analysing mortality risk in insurance-linked securities – in particular, life settlements. Traditional underwriting does not work well at older ages and with the segments of population most likely to sell their insurance policies to investors. Medical experts play a major role in properly estimating life expectancies in the life settlement context, especially for individuals with impaired health. On the other hand, technical or actuarial analysis is, of course, the foundation of building any methodology for estimating LEs. Integrating the two is a common challenge faced by LE providers.

Knockout versus debit/credit approach

The two main approaches to determining an individual's health condition and corresponding mortality rates are the knockout approach and the debit/credit system. Under the knockout approach, also referred to as the edge approach, not meeting a specific requirement leads to placing the individuals in a higher risk class with high mortality rates. Specific underwriting guidelines determine the criteria that need to be met. It is the most common method in preferred underwriting. In the debit/credit approach, debits and credits are determined for each of the underwriting criteria, and their sum determines the placement of the individual in a particular risk class. It is a point system in which points, either positive or negative, are assigned for each of the underwriting criteria.

DEBITS

A system of debits and credits is used in the life insurance industry to determine the loading to a "standard" mortality due to health impairments or other factors discovered in the underwriting process. For example, a cardiovascular disease calls for a substantial debit, as it indicates a sizable increase in mortality rates relative to individuals who do not have any cardiovascular problems. Numerous factors are used in underwriting, and each of them could in theory be assigned a certain level of debit or credit based on

how it is expected to affect the mortality rate. The sum of all debits and credits, with debits considered positive and credits negative, is then the loading applied to the mortality table designated as standard. The loading is typically expressed as a percentage of the standard mortality rates, with debits of 175 being equal to a 275 “table rating”. That means that the probability of the insured’s dying the next year is 2.75 times that of a standard life. A table rating of 100 corresponds to the standard mortality table with no adjustments. If D is the total debits, the mortality multiplier is $1 + D$. In other words, the adjusted mortality rate q'_x is determined as $q'_x = q_x (1 + D)$, where q_x is the mortality rate for age x taken from the standard mortality table.

In the simplest approach, the same mortality multiplier is applied to mortality rates at all ages, and not only to the mortality rate in the first year after underwriting. The probability of the individual underwritten at age x dying between the ages of $x + t$ and $x + t + 1$ is then determined as $q'_{x+t} = q_{x+t} (1 + D)$.

In the life settlement context, this approach when applied blindly leads to some illogical results. Since life settlement transactions typically involve older lives with somewhat impaired health, mortality rate q_x is large to begin with. In addition, debits reflected in the mortality multiplier could be very high, often exceeding the level of insurability if the individual were to attempt to purchase a new life insurance policy. Mortality rates loaded for debits could relatively quickly reach 100%, leading to the conclusion that after a certain age the individual is certain to be dead. This is rarely the case in reality. However, it is the mistake still made by those LE providers who blindly apply a constant multiplier factor to an old-age mortality table. The resulting shape of the mortality curve in many such cases is completely unrealistic.

There are several solutions to this problem that could be implemented within the same general framework. One of them includes changing the multiplier factor as opposed to keeping it constant. For example, the multiplier could be reduced each year, similar to the way select mortality tables trend towards the ultimate rates as the effects of underwriting wear off. The multiplier could also remain constant for a number of years and start reducing after that.

One of the other approaches sometimes used involves looking at the survival probability p_x as opposed to its complement, the rate of mortality q_x . It is then assumed that the survival probability for an individual is the standard survival probability to the power equal to what has previously

been referred to as the mortality multiplier, that is, $1 + D$. That is, $p'_x = p_x^{1+D}$. Written in terms of mortality, this becomes the following

$$1 - q'_x = (1 - q_x)^{1+D}$$

From here we can write that the substandard mortality rate $q'_x = 1 - (1 - q_x)^{1+D}$. Using binomial series expansion, we can then write

$$q'_x = 1 - \sum_{n=0}^{\infty} \binom{D+1}{n} (-q_x)^n$$

This could also be written in the following form

$$q'_x = (1 + D)q_x - \frac{(1 + D)D}{2!}q_x^2 + \frac{(1 + D)D(D - 1)}{3!}q_x^3 - \dots$$

Assuming smaller values of mortality rates and debits, we can approximate q'_x by taking the first term of the series, leading to the familiar equation $q'_x = q_x(1 + D)$. For older ages and greater mortality impairments typically found in life settlements, this assumption does not hold well, since the mortality rate could reach 100%, resulting in an unrealistic shape of the mortality curve as mentioned above.

The definition of the debits and credits has effectively been changed, and the values determined in the traditional way might not be applicable any more.

As to whether credits and debits are fully additive or should be considered multiplicative, the answer is that it should really be the combination of the two. Practical implementation of this approach, however, is not always feasible.

CHOICE OF MORTALITY TABLE

The choice of a reference mortality table is critical in estimating LE and general mortality characteristics. The mortality rates resulting from the analysis might be very different from the ones in the reference mortality table; such a discrepancy does not detract from the importance of having a set of tables serving as a reference in analysing an individual's mortality characteristics. Even those LE providers that claim to have developed their own mortality tables have begun with a set of established traditional mortality tables.

Investors and LE providers are presented with a very difficult choice in deciding which mortality table to use in their analysis. There are a number of mortality tables available to them. None of them is directly applicable to the life settlement populations; all have been constructed for purposes very

different from that of assessing life settlement mortality. Some of the LE providers have accumulated experience data on life settlement mortality, but this data has limited statistical credibility. The underwriting criteria used by most LE providers have been changing over time. It is difficult to bring all experience data to the common denominator to make reasonable comparisons. A strong argument can be made that the 2008 Valuation Basic Table (2008 VBT) described below should be used in constructing life-settlement-specific mortality tables, even though it is undeniable that significant changes to the table have to be made to reflect life-settlement mortality.

Understanding how a particular table has been constructed is important for its proper use, particularly when a mortality table is used for a purpose different from the one for which it has been developed, and when it is applied to a population different from the one that produced statistical data used in the table construction. Furthermore, knowing the assumptions used in the construction of a mortality table, as well as its specific limitations, is necessary if the table is to be modified to make it applicable to mortality-linked securities such as life settlements.

2008 VALUATION BASIC TABLE

The latest mortality table for insured lives produced in the US is the 2008 Valuation Basic Table, commonly referred to as 2008 VBT. It was developed to replace and improve on the 2001 VBT table by incorporating new statistical data and addressing some of the specific weaknesses of the 2001 VBT.

Even though the 2008 VBT is intended for valuation purposes in life insurance, it does not incorporate the margins included in the valuation process by life insurance companies. In this sense it is similar to the 2001 VBT. It is worth noting that even the older 2001 VBT has not become the standard or the starting point in pricing for many insurance companies. Some insurance companies are using even older tables while making their own adjustments to them; some use their own experience and largely disregard the VBTs for pricing purposes. Since risk classification and underwriting criteria differ from one company to another, the use of individual mortality tables is understandable. The question of whether the shape of the mortality curve is appropriate remains open. There is strong opinion that a reasonable approach would be to use the 2008 VBT as the foundation for building custom mortality tables based on specific experience.

Chapter 13 mortality and longevity models described some of the basic concepts involved in mortality table construction. This chapter builds on that foundation and introduces some specific ways to construct a mortality

PANEL 14.1 PROJECTION PURSUIT REGRESSION

Projection pursuit regression is one of the predictive modelling techniques that are sometimes employed in fitting a mortality table to the available data. The Society of Actuaries team working on the development of the 2008 mortality VBT utilised projection pursuit regression along with the Whittaker–Henderson method for this purpose. It has been argued that the choice of projection pursuit regression is far from optimal for the task, and that better methods are available. The resulting table had such a significant element of judgement involved at many steps of its development, that the choice of projection pursuit regression over other methods as the predictive modelling technique probably has had little impact on the end result.

Projection pursuit regression is one of the generalised additive models. It does not belong to the category of generalised linear models often used in the analysis of casualty insurance risk. It represented a significant step forward when initially developed for applications in high-energy physics, but projection pursuit regression is rarely used nowadays and is usually replaced by approaches based on neural networks for the more demanding applications, and by the simpler generalised linear models whenever possible. Numerous other approaches could also be utilised. Projection pursuit regression could be seen as a rather general approach including, for example, a simple linear regression or a neural network with one hidden layer.

In a standard regression model, $Y = f(X) + \epsilon$, where Y is the observation based on the predictor X . More precisely, X is a multidimensional explanatory vector and Y is the response variable, with X and Y in the model forming an observable pair of random variables from a distribution, while ϵ is the error term independent of X . The aim of regression analysis is to estimate the conditional expectation of the response variable given the explanatory variable, or $f(x) = E[Y|X = x]$. The standard projection pursuit regression approach involves approximating $f(x)$ by a finite sum of the so-called ridge functions that are in turn different linear combinations of X_k . If the size of the available random sample of the explanatory vector and the response variable is n , ridge functions are defined on dimensional spaces significantly lower than n (in practical applications usually very low). For each direction, or projection pursuit, the process of fitting is effectively a univariate smoothing. The iterative approach calculates the projections and the sums of the resulting ridge functions to improve the model's fit to the

data. The process uses residuals from previous “smooths” when repeated for another direction. In constructing the 2008 VBT mortality table, fit was determined by the straightforward approach of minimising the sum of squared errors between the residuals and the sum of the next iteration ridge functions. This is the most common way that could be easily implemented using available software tools.

Projection pursuit regression is a nonparametric procedure that does not impose any specific relationship or constraints on the relationship between the explanatory vector and the response variable. It then goes further by overcoming the well-known limitations of most nonparametric regression methods, since it does not utilise recursive partitioning that could lead to unnecessarily complex and difficult-to-interpret models. Easy graphical interpretation of results is another advantage of the method. Projection pursuit regression also has some known disadvantages, such as the potential for oversmoothing, and the difficulty it has in modelling regression surfaces that vary almost equally strongly for all possible linear combinations used by the model.

table based on the available historical data. Panel 14.1 describes one of the predictive modelling techniques used in developing the 2008 VBT.

A significant element of the 2008 VBT is the incorporation of population mortality data in constructing the table. Such information sources considered in the construction of the 2008 VBT include the Social Security Administration data based on Medicare death records, data from the Centers for Disease Control (also based at least in part on Medicare records), and Veterans Administration data – especially for older ages, where the data is limited and lacks statistical credibility. These are the ages of most relevance to investments in mortality risk such as life settlements. It is notable that for the first time the table is extended to the age of 120. The omega mortality rate for ages 110 and older was chosen to be 0.45, based in part on research into older-age mortality.

In constructing a mortality table, it is important to find the right balance between fit and smoothness. Panel 14.2 provides a brief description of the Whittaker–Henderson method used for that purpose in constructing the 2008 VBT. It is a relatively common way to fit and smooth mortality rates to large data samples.

Extending the Whittaker–Henderson method into any area with limited data, of which older ages is the most obvious, could produce unreasonable

PANEL 14.2 WHITTAKER–HENDERSON METHOD

The Whittaker–Henderson graduation method was used in the development of the 2008 mortality VBT as well as the previous 2001 mortality VBT. The method involves determining graduated mortality rates by minimising the function $F + c S$. F and S are defined the following way

$$F = \sum_{i=1}^n w_i (q_i^{\text{graduated}} - q_i^{\text{raw}})^2$$

$$S = \sum_{i=1}^{n-r} (\Delta^r q_i^{\text{graduated}})^2$$

q_i^{raw} is the raw (not graduated) rate; $q_i^{\text{graduated}}$ is the graduated rate; w_i is the weight or exposure assigned to the q_i^{raw} value. Δ^r is the difference operator of r th order, that is

$$\Delta^r q_i = \sum_{k=0}^r (-1)^{r-k} \binom{r}{k} q_{i+k}$$

F represents a measure of fit, while S represents a measure of smoothness. Smoothness is measured for graduated rates as opposed to the raw rates. The constant c is effectively used to obtain a weighted average of the two. A low value of c emphasises fit over smoothness; a high value of c puts more emphasis on smoothness over fit. It is instructive to note that in the extreme case of zero c the procedure returns raw data.

results. The weights used in the fitting procedure will be low, while the smoothing procedure uses equal weights; the result will be oversmoothing and a poor fit in these areas. In addition, using a third- or higher-order procedure could artificially overstate mortality rates at the extremely old ages. One has to be mindful of these limitations in using the table.

While numerous criticisms of some elements of the 2008 VBT have already surfaced, there is a general recognition that this mortality table represents a significant advance and has advantages over other industry tables in most applications related to insurance-linked securities.

RELATIVE RISK RATIOS

The 2008 VBT uses the concept of relative risk (RR). RR tables correspond to various risk classes with respective mortality levels; they are based on the Primary Tables. RR tables represent mortality levels for specific segments of the insured population; they correspond to the relative risk ratios (RRRs) used to determine an individual's mortality risk level.

RRRs differ by age, gender and smoker status. For each issue age, gender and smoker status, RRRs are multiplier-type factors that determine the relationship between the preferred and the aggregate mortality. The overall aggregate mortality for fully underwritten lives is assigned the relative risk of 100%. The convention is to refer to such a table as RR100. The RR70 table corresponds to the 70% relative risk. It is important to note that, even though the weighted average mortality for RR70 is 70% of the RR100 table, the multiplier for a specific subset will likely differ from 0.7, as the difference between the two tables would not necessarily be 0.7. For example, the adjusted multiplier for issue age 25 male non-smokers is 0.8, and for issue age 65 male non-smokers it is 0.65 (in both cases comparing the RR70 and the aggregate RR100 tables). Significant judgement went into determining the preferred wear-off factors for the RR tables.

There is a direct relationship between the RRRs and another measure of relative risk, the underwriting criteria score (USC). The USC can be used to assign a specific mortality table to a given risk. For each age, gender and smoker status the relationship between USC and RRR differs. The Society of Actuaries developed a conversion table to determine RRRs corresponding to specific USCs. For practical uses, the Expanded Conversion Table should be referred to. The conversion algorithm is relatively simple. Standardised underwriting criteria were developed to assign underwriting criteria scores to individuals, thus avoiding the problem of two insurance companies using slightly different underwriting criteria. They would do this usually by using “exceptions” to place an individual in a specific risk class, even though he would technically belong in a riskier class based on the rather inflexible rules of the knockout rating approach.

In determining mortality levels for an individual, the USC is calculated based on the standard underwriting criteria. It is then converted to an RRR, which in turn determines which valuation mortality table is appropriate.

UNDERWRITING FOR OLDER AGES

Life insurance underwriting is the process of assessing medical and other data about an individual and determining their risk classification and mortality rates, which in turn determine LE.

Medical underwriting in life settlements is focused on the older ages, as these dominate the life settlement market. The underwriting process greatly depends on the LE provider; there is little uniformity.

In most cases, reinsurance underwriting manuals such as the one used by Swiss Re are taken as the initial basis for medical underwriting. LE

providers could then introduce their own modifications to the manual and build on them. The use of reinsurance manuals introduces some degree of standardisation, without which it is impossible for a reinsurer to compare data and perform its underwriting.

Another reason for using reinsurance methodologies is that reinsurers have the most expertise in old-age underwriting. Policies with large face values are most often taken out by individuals over 45 years old. Even though the age range from 40 to 60 – the most common for large face-value policies – is well under the 65+ category in life settlements, it is the closest available. Moving to older ages such as 75 and 85 requires additional adjustments.

We have discussed some of the issues in applying the standard debit and credit approach to older ages and substandard mortality classes. The debit and credit system was never intended to be used for such cases, and the high debits common in substandard classes in older ages would not be used in traditional life insurance underwriting. For example, total debits of over 500 typically make a person uninsurable, and insurance companies do not analyse mortality for such individuals. Individuals in this category would usually not be issued a life insurance policy, and there is no mortality table or guidance applicable to them in the underwriting manuals utilised by the insurance industry. While it is technically possible to calculate the total debit D and apply the $(1 + D)$ multiplier to a mortality table, the results are unreliable and often inconsistent. However, even though this approach is obviously wrong, it is sometimes useful in obtaining a reference point to compare with the results of other methods. Some of the possible adjustments to make this approach reasonable, at least to some degree, have been described above.

The distortion introduced by employing the debit and credit approach to substandard older lives¹ could be amplified by applying the resultant debits to a wrong mortality table. The choices available in selecting a mortality table have also been discussed, with a clear conclusion that there is no good choice, even though some are better than others. The degree of distortion resulting in using a debit-based mortality multiplier and mortality rate taken from an inappropriate mortality table is quite remarkable and is not fully appreciated by most of the end-users of the LEs, most importantly by investors in life settlement paper. The most important indicator of the distortion is the completely unrealistic shape of the mortality curve often produced by this approach.

As experts in older-age mortality, LE providers have made significant

adjustments to the standard underwriting manuals. The adjustments are based on the data collected on the actual mortality relative to the previous determinations of LE. Unfortunately, this data sample is still limited in size due to the infancy of the life settlement industry. For the same reason, the data is skewed toward shorter LEs and mortality in the first years after LE assessment. While the credibility of the data is limited, it could still be used to make some of the adjustments. Additional adjustments are based on the research of longevity associated with specific diseases of the elderly, as new advances in treatment are analysed and incorporated in the analysis.

Assessment of longevity for the most serious diseases, such as specific types of cancer, is an area where significant value could be added. In such cases, clinical judgement could override other considerations and any standard mortality tables. Some of the LE providers have performed their own analysis of available data or have partnered with other organisations to come up with longevity projections for specific diseases. An example would be the analysis of mortality of people with Alzheimer's disease at various stages, taking into account co-morbidity with other medical conditions.

Standard underwriting tools

Below are some of the standard life insurance underwriting tools used in addition to the information contained in the life insurance application. The application information would typically not be present in life settlement underwriting.

- ❑ An Attending Physician's Statement (APS) is a standard underwriting tool used in life insurance regardless of the applicant's age. It provides valuable information about the applicant's health and medical history. Where warranted, it is supplemented by statements from specialists who might have observed or treated the applicant.
- ❑ A blood profile is used to screen the applicant for medical conditions such as diabetes, HIV, kidney and liver diseases, potential of cardiovascular disorders and many others. In life settlement underwriting, a blood profile might not be available or the underwriter might have to rely on the results obtained from the applicant's physician. These results might not be recent.
- ❑ Urinalysis could provide an indication of kidney diseases, certain tumours and other medical conditions. It could also verify the applicant's non-smoker status as well as drug usage. As with the case of blood work, it might not be available in life settlement underwriting.

- ☐ EKG and stress tests could be used to screen for or assess the severity of cardiovascular conditions.

The standard categories are of equal importance in traditional life insurance underwriting and in underwriting for older ages in the life settlement or other context. These include:

- ☐ age;
- ☐ sex; and
- ☐ smoker status.

Other useful categories in life settlement underwriting include marital status, policy face amount, geographic location, family medical history, and income level. Any detailed information on the above could add value. For example, elderly men have an increased mortality rate after the death of their spouses, which after a certain period of time reverts to the mortality rates expected otherwise. Since most LE providers also use their judgement in underwriting, any additional information could add value to the underwriting process.

Additional underwriting tools

Important elements of the older-age underwriting process that are not always used in traditional life insurance underwriting include the following.

- ☐ Assessment of daily activities could provide an indication of the lifestyle, which, everything else being equal, often correlates with LE. Commitment to a healthy lifestyle, to the degree it could be assessed by the underwriter, is a potential input in the underwriting process. Access to caregivers and a support network fall in the same category.
- ☐ Cognitive impairment testing could be performed over the phone and is an effective way to screen for Alzheimer's and similar conditions. It consists of a number of questions designed to assess an individual's mental alertness and memory.
- ☐ Credit score and a consumer credit report are sometimes used in life settlement underwriting. In traditional life insurance underwriting, their direct use is not always allowed. Even if they are not used in life insurance underwriting for the purpose of assigning risk category, determining insurability or setting premium levels, they could be used as one of the tools for identifying applications to be further reviewed to ensure no STOLI issues are present.

Another consideration in life settlements is that the insured have a financial incentive to appear sicker, since that increases the value of their policies. There have been reports of life insurance agents actively encouraging their clients to visit specialists with complaints before their LEs are assessed. This calls for caution in applying life insurance underwriting techniques to life settlements, since life insurance applicants have the opposite incentive and want to appear healthier in order to be assigned a lower risk rating.

Additivity of debits and credits

The system of debits and credits in life insurance underwriting was designed in such a way as to ensure that debits and credits are additive. When individual debits and credits are small, adding up all debits and credits to calculate the mortality multiplier produces a result not too different from calculating multipliers for each of the debits and credits, and then multiplying the multipliers to compute the total mortality multiplier. When the magnitude of the debits grows – as in older-age underwriting in general, and in particular for those with greater health impairments common in life settlements – this difference becomes quite significant. It is the difference between $1 + \sum D_i$ and $\prod (1 + D_i)$. Obviously, the definition of D_i might be different in both cases. An argument has been made that considering debits as additive could introduce a distortion leading to understating life expectancies, and that the multiplicative approach avoids this distortion. More complex approaches have also been suggested – in particular, adding debits and credits within main categories, and then multiplying the mortality multipliers corresponding to each of these categories. Another suggested approach treats debits as truly multiplicative instead of as additive. Or, as previously mentioned, we can calculate partial multipliers and then treat the product of these multipliers as the mortality multiplier. Without modifications, this method is inconsistent for values of D_i that are less than 100%. Again, it is important to point out that the change in the way debits and credits are used is, effectively, a change in the definition of the debits and credits and the way they are determined.

This question becomes very important in the life settlement context because of the sheer size of the debits and the magnitude of the adjustment to mortality rates. While the additivity of debits and credits is sometimes discussed, other approaches are not known to be widely used.

CHOOSING THE LIFE EXPECTANCY

In practice, choosing the LE in the broader sense of choosing the mortality rates for an individual is almost equivalent to choosing an LE provider. LE is the biggest uncertainty and risk in life settlements. An investor is typically shown LE reports from more than one LE provider and is presented with a difficult choice in deciding which of them is more accurate. The range of LEs for the same individual could be very wide.

The groundhog is like most other prophets: it delivers its prediction and then disappears.

Bill Vaughan

All of the parties involved in a life-settlement transaction, from the insured to the brokers to the LE providers, are paid before, at or shortly after the closing of the transaction. After that, their involvement is limited since they have received financial compensation for their role in the transaction. The exception is the investor who has paid for the life settlement and taken over the cost of paying premiums and expenses, in the expectation that these negative cashflows will be followed by a positive one in the form of the policy benefit. If this positive cashflow comes later than expected, the investor realises a lower-than-expected rate of return and might even suffer a loss. Every party to the transaction has a relative certainty as to the financial result of the life settlement transaction; the investor is the only one left with the uncertainty to be resolved years in the future.

LE providers are seen as having their short-term financial interests misaligned with those of investors. LE providers are usually engaged and paid not by investors but by brokers or life settlement providers. Brokers and providers are interested in a speedy consummation of the transaction. In some cases, they also have a financial interest in getting the highest price for the policy. The way to increase the price is to make the policy appear more valuable by understating LE.

The existence of perverse incentives does not imply that any of the LE providers is dishonest or unprofessional. LE providers strive to develop the best methodologies and the most accurate ways of estimating life expectancies. Furthermore, the perverse incentives are a factor only in the short term; over a period of several years any understatement of LEs would become clear and damage the business franchise of the LE provider.

LE providers try to provide the most accurate estimates to investors, but their results could still differ by a significant degree for a specific insured, even if they produce the same average LEs and there is no general bias. Investors should develop some internal expertise in this area, since

mortality is the single most important determinant of the investment performance of life settlements and related securities.

LIFE EXPECTANCY SHOPPING

The price an investor is willing to pay for a life insurance policy varies inversely with LE. There are no established methodologies for determining individual mortality rates and LE except for those used by insurance companies in underwriting life insurance applicants. LE providers used by investors often produce widely varying estimates. In most cases, these estimates are ordered by brokers and not by the ultimate investors; the brokers are then presented with more than one LE estimate. A broker might sometimes present to investors not all the estimates but only those indicating higher mortality rates and lower life expectancies. This introduces a bias in the analysis performed by the investor based on the provided information.

It is also known that some LE providers produce lower life expectancies for specific health impairments or general risk categories. Some LE providers are even believed to produce lower life expectancies than others for all risk categories. This allows a broker with knowledge of such biases to approach only those LE providers believed more likely to provide shorter life expectancies in a particular case, and to avoid those more likely to come up with lower mortality rates. The broker is then providing investors with all the LEs obtained; so, while none are hidden or omitted, the bias is still there. This phenomenon is sometimes referred to as LE shopping.

Some investors have found themselves defenceless against this tactic and have ended up accumulating portfolios of life settlements based on understated life expectancies. This improper portfolio valuation due to hidden bias can go unnoticed for many years.

For investors, one way of avoiding this risk involves establishing a closer relationship with LE providers, becoming comfortable with their methodologies and insisting on obtaining life expectancies from those particular LE providers. More sophisticated investors go beyond these steps and develop an internal view of the biases of individual LE providers, seeking to better understand the methodologies used. Based on this understanding, investors can make their own adjustments to the analysis performed by an individual LE provider, or can decide that the results produced by some LE providers lack any credibility and do not add value.

Additional bias

Brokers trying to present shorter LEs to investors could introduce a systematic bias in one more way. They can encourage policyholders interested in

settling their policies to undergo a general physical checkup, to visit medical specialists and to fully state their health complaints. As the medical file grows, the LE providers are more likely to see evidence of health problems and come up with shorter LE estimates. Even if no new medical problems have been discovered, the very fact of more visits to a doctor could lead to higher mortality rates produced by LE providers.

In this case, there is no suggestion of any impropriety, as the medical complaints are real; there is no false information produced or any facts omitted. However, the strategy leads to the same effect of understating LEs: mortality rates are assigned based on comparing an insured to some average for which a mortality table has been developed; and for this average person, there have been no additional medical visits made or health complaints reported at the suggestion of a broker. The difference could be even more pronounced when the basis of mortality rates used by an LE provider is a mortality rate developed by the insurance industry based on life insurance underwriting. In life insurance underwriting, the tendency is to underreport certain health conditions to qualify for lower rates. This bias is reflected in the mortality tables based on statistical data on mortality experience of insurance companies. The opposite is true in life settlements. The more sophisticated investors have learned to identify brokers and other sources of business that employ this technique and make specific adjustments to account for the potential bias.

ASSUMED PREMIUMS

When pricing life settlement securities we must consider the amount of premiums required to keep policies in force. In the case of Universal Life insurance, many unsophisticated investors have made the mistake of believing the premiums will always be based on the current assumptions used by the insurance company, most importantly, the current declared net crediting rate. Premium levels provided in the policy illustrations are not guaranteed. Often, they can be increased by the insurance carrier in the future because of the lower net crediting rates, higher-than-expected mortality, or other reasons.

How long investors have to wait to realise their returns has an effect on the return level. The longer they have to wait, the lower the rate of return. This is more than a question of the time value of money. As policyholders live longer than assumed in the projections, the investors have to continue paying premiums to maintain the policies in force.

BEING PAID FOR THE RISK

An investor in a mortality-linked security is exposed to certain risks and wants the investment return to provide compensation for assuming these risks. In this sense, it is similar to investing in any security.

The main risk involved in life settlements is the mortality risk. In pricing a financial instrument such as an insurance-linked security, the price depends on the cashflows associated with the security. Cashflows associated with a life settlement security could be highly uncertain due to the fact that mortality rates, even when determined precisely, are probabilities and not deterministic measures. In addition, there is significant parameter risk involved, due to the fact that mortality rates can never be determined precisely and might even have a consistent bias.

The uncertainty as to the amounts and timing of cashflow affects the yield an investor would require when investing in such a security. The uncertain and possibly quite long time horizon prevents many investors from participating in this market at all.

In practice, the discussion of the theoretically correct prices is purely academic, since the market for these securities is highly inefficient. Many, if not most, of these securities are mispriced by almost any measure, but this creates unique opportunities for the more sophisticated investors operating in this space.

It is noteworthy that there is some optionality present in investing in such securities. An investor might re-evaluate the mortality and other assumptions and decide to stop paying premiums and drop the policy.

Investment risks unrelated to mortality

Many of the risks of investing in these securities are unrelated to mortality. Other chapters provide a more comprehensive description of these risks and their effect on the required yields. The risks are best assessed and managed in the context of portfolio management, explored further in the following chapters. Under certain circumstances, these risks might even be the primary drivers of investment performance and price, as opposed to the price being determined mostly on the basis of mortality characteristics.

CONCLUSION

This chapter builds on the concepts and practical aspects of mortality modelling that have been introduced earlier, from the point of view of the investor in mortality and longevity risk. While the focus here is on life settle-

ments, many of the same concepts could be applied to the analysis of other types of mortality-linked securities.

LE is the most important input in a model for valuing life settlements. To perform more than rudimentary analysis, mortality rates for an insured should be estimated beyond simply one number for LE. So-called LE providers supply such data as an output of their analysis.

The 2008 Valuation Basic Table is becoming a standard reference point for all mortality calculations involved in estimating LE of individuals. The table is a clear improvement over its predecessors but has its own limitations. The introduction of a truly standardised underwriting criteria score (USC), along with the relative risk tables, represents an important advance, which is particularly relevant in mortality-linked securities such as life settlements.

Overreliance on LEs and mortality estimates obtained from third parties could lead to substantial valuation errors. An investor has to develop some expertise in this asset class before venturing down the road of investing in mortality- and longevity-linked securities. The model of fully relying on third parties for policy valuation and all other services has become discredited and its dangers revealed. The need for developing in-house expertise is becoming more apparent to investors, even though the process is very slow.

Once a policy enters an investment portfolio, there is usually little concern with risks other than mortality, and little effort goes into their quantification. It is often assumed that, if a policy has been purchased, sufficient due diligence has been performed at the point of purchase and there is no need to further analyse and monitor these exposures. In fact, however, legal and other risks not directly related to mortality performance might have a great impact on the portfolio valuation and realised investment returns. These factors have to be carefully analysed and taken into account in the modelling process by any investor.

Distribution of deaths is unknown in a portfolio of life settlement securities. The mistake made by many a naïve investor has been in assuming that the cashflows are highly predictable; this is true only in the case of very large investment portfolios, and where there are no biases in the mortality estimates. It also requires the assumption that other types of risk are minimal, which is not always true for this asset class. Mortality is the single most important driver of investment performance, but it is not the only one.

Portfolio analysis is particularly important in the context of investing in life settlements. The following chapters provide an overview of portfolio management, the risks involved, and the types of analysis that could be employed where life settlements and other mortality-linked securities are

concerned. The emphasis is on stochastic approaches in modelling in order to properly manage risk and to maximise risk-adjusted return.

- 1 Terms such as “substandard lives” in reference to individuals with impaired health are accepted in the life insurance industry and do not carry any negative connotation.

Longevity Risk Transfer and Longevity-Linked Securities

This chapter analyses longevity risk, and looks at ways to transfer it to the capital markets and to invest in it. It examines the effect of potential longevity improvements on defined benefit (DB) pension plans as well as the implications of such improvements for annuity providers and other holders of longevity risk. Insurance-linked securities such as longevity derivatives and longevity bonds are examined from the point of view of both investors in and hedgers of longevity risk. Some of the actual transactions in the UK, where most activity in longevity risk transfer has so far taken place, are also described. Such transactions – from pension plan buyouts and longevity insurance offered by Pension Corporation, to the £500 million longevity swap between Canada Life and JP Morgan, to the longevity risk transfer of £3 billion in BMW's UK pension liabilities to Deutsche Bank – demonstrate the potential for future growth of the longevity risk transfer market. In addition, special attention is paid to the unique longevity risks faced by investors in traded insurance policies.

LONGEVITY RISK

Longevity is so closely linked to mortality that one term is usually defined in relation to the other. Longevity risk has been overlooked or underestimated for a long time, and it still is, despite the growing recognition of its significance.

Definition of longevity risk

Longevity is simply the opposite of mortality. As defined in Chapter 13, longevity is the probability distribution of an individual's staying alive over a certain period of time or beyond a certain age or point in time. Greater longevity corresponds to lower mortality, and vice versa. The concept of longevity is generally applied to populations rather than individuals, even

though the calculations are based on longevity of individual members of a population.

In addition to viewing longevity in a probabilistic framework, in largely the same way as is done for mortality, one can speak about actual or realised longevity of a population or individual.

The term “longevity risk” is used when greater-than-expected longevity (lower mortality) has detrimental financial effects; the term “mortality risk” is used to describe the possibility of a negative financial impact of greater-than-expected mortality (lower longevity). Longevity risk is typically considered over a long time horizon, while mortality risk can refer to both long and short time periods. An example of the latter is a mortality spike, the risk of which can be transferred from insurance companies to the capital markets in the form of extreme mortality bonds (described in Chapter 11).

Entities and securities exposed to longevity risk

A clear example of a longevity risk holder is an insurance company selling annuity products¹ for which higher-than-expected longevity of annuitants results in lower profits. The same type of longevity risk is present in DB pension plans, where underestimating the longevity of plan participants means that payments will need to be made over a longer period than assumed, creating a potential unfunded liability. The biggest holder of longevity risk is usually not private pension plans but governments. (In the case of governments, as significant as the longevity risk could be, it often pales in comparison with much greater issues of government pension or social security systems being completely or partially unfunded in many countries. Given the demographic shifts in the developed countries, the pay-as-you-go system adopted by these governments is not sustainable without the introduction of major changes.)

Longevity risk is also a major factor in portfolios of life settlements (traded life insurance policies). Longer-than-expected longevity leads to lower-than-expected returns on life settlement portfolios and, in extreme cases, to portfolios having negative net present value (NPV). In life settlements, longevity risk is considered in reference to a very small subset of the general population or even of an insured population. Longevity is the main investment risk in most life settlement portfolios.

Reverse mortgages² present another example of longevity risk. One more example involves health insurance. In some cases, in countries where medical insurance is not free or is free only at a certain basic level, employers have plans under which retirees receive free or subsidised medical insur-

ance coverage. Often, the pension plans that provide DB payments to these retirees provide the medical coverage. Technically, the pension plans provide the pension payments but pay a third party (a health insurer) to provide medical benefits. It is also possible that the medical benefits are provided by a defined contribution (DC) rather than by a DB plan. Terminology can differ: technically these plans might be called payees rather than providers, when the term “provider” is reserved for the insurance company administering or actually providing medical coverage. The longer the plan participants live, the longer such benefits have to be provided, creating exposure to longevity risk.

Leaving aside the longevity risk borne by central governments – since it generally cannot be transferred to investors – the main concentration of longevity risk appears in the following areas:

- ☐ private defined benefit pension plans (referred to as pension schemes in the UK) or plans run by local governments;
- ☐ annuity providers;
- ☐ life settlement investors;
- ☐ holders of reverse mortgages;
- ☐ providers of free or subsidised medical insurance coverage to retirees; and
- ☐ investors in the entities and insurance-linked securities that are exposed to longevity risk.

NEED TO TRANSFER LONGEVITY RISK

When actual (realised) longevity is longer than expected it can lead to funding shortfalls and the emergence of sizable unfunded liabilities for a pension plan or annuity provider. In the case of pension plans – as the largest area where longevity risk resides – relatively small increases in longevity of the pension plan participants can lead to significant increases in pension plan liabilities. The calculation of liabilities is typically based on assumptions prescribed by the government directly regulating pension plans, or by the government tax authorities that establish a set of conditions for pension plans to maintain a favourable tax status. These assumptions include primarily the investment assumptions (such as discount rate) and the selection of appropriate mortality tables. The prescribed investment assumptions can include a certain safety margin (though often they are criticised as overly optimistic in the current investment environment); application of a prescribed mortality table, however, results in a determin-

istic outcome that fails to take into account the variability of actual results around the expected mean – that is, longevity risk. In fact, the issue of longevity goes even deeper and concerns the very choice as well as the interpretation of mortality tables. If mortality tables show greater mortality rates than is appropriate for the specific population of the pension plan participants, the likely result is a shortfall and unfunded pension liabilities. The true longevity risk manifests as a decrease in mortality rates that has not yet been observed but can very well occur in the future, making the current mortality tables inapplicable and resulting in pension plan liabilities being understated.

The risk of longevity improvements is very real and is based on both the observed trends and potential future changes of the factors that affect longevity. The degree of longevity improvements is difficult to predict, creating the need to hedge this risk to preserve the ability of pension plans to meet their obligations to the plan participants.

To take another example, that of investing in traded life insurance policies (life settlements), the risk of greater-than-expected longevity is tied less to the overall longevity improvements of the general population and more to the longevity of the specific population of the life settlements investment portfolio. Longevity and its changes for the insured lives in such a portfolio maintain a degree of dependency on the longevity and its changes for the general population; but the more narrowly defined population of the insured who have chosen to settle their policies (see Chapters 12, 13 and 14) tends to differ significantly in its mortality and longevity characteristics from both the general population and the insured population. While the risk of longevity improvements is a factor, an even more important risk comes from the longevity having been significantly underestimated from the very beginning, irrespective of any potential longevity improvements in the future. The risk of longevity being greater than expected (labelled “extension risk” in the life settlements vernacular) is critical in the investment analysis of life settlements; as greater-than-expected longevity has resulted in many significant losses that could have been prevented by a more accurate evaluation of the risk. Most of these losses have been the product of poor initial analysis and a systematic understatement of life expectancies (LEs). These are known issues concerning systematic biases in the way life settlements used to be (and sometimes still are) priced. Going forward, however – especially with more life settlements that have long LEs – the longevity risks having to do with potential longevity improvements, and with the longevity of a specific pool of life settlements being greater than

expected due to statistical fluctuations, need to be hedged or addressed in some other way. Some of the same longevity risk transfer instruments that can be used for pension plans are applicable to life settlements as well (though they would be used in a different way and would utilise different longevity indexes or reference pools).

To summarise, there are three main types of longevity risk that can have a significant financial effect on the holders of the risk:

- ❑ longevity improvements that reflect changes in the overall mortality rates resulting from the trend of a population living longer;
- ❑ random statistical fluctuations around expected longevity (fluctuations around the mean) that occur even when the mean does not change due to longevity improvements, and in particular when the sample size is not sufficiently large; and
- ❑ underestimating the true longevity of a population, by applying a wrong set of mortality tables, using inappropriate mortality estimates, or due to other modelling errors, irrespective of the risk of unanticipated longevity improvements.

The last type stops being a risk once the mistake of the initial underestimating of the mean longevity becomes obvious, but it is a risk until that happens. In addition, even when the mistake becomes obvious or appears to be likely, the market inefficiency may not make it obvious to other market participants and may allow hedging, or transferring the risk below cost. This statement refers primarily to the life settlements market, which suffers from inefficiency, rather than to pensions or annuities. It may also be that two parties simply have different opinions on the longevity of a certain population, and are willing to enter into a transaction fully aware of the opinion of the other party.

To illustrate the magnitude of longevity risk facing defined benefit pension plans, Pension Corporation uses the assumption of one year of longevity extension corresponding to the 3.5% increase in pension plan liabilities in the UK.³ (The 3.5% figure is representative of the general sensitivity of pension liabilities to longevity improvements in most of the developed countries.) This prominent example clearly shows that longevity is a major risk that cannot be neglected.


The need to address the risk of longevity is critical for pension plans, investors in life settlements, annuity writers and other parties that may not even realise the degree to which they are exposed to longevity risk. This

degree can be quite significant and is increasingly being recognised by holders of longevity risk who would like to offload or minimise their exposure.

LONGEVITY IMPROVEMENTS

Lifespan extension has been observed in most countries over recent decades. Chapter 13 demonstrates these longevity improvements in the US and UK based on government statistical data.

The two primary reasons for longevity improvements are better living conditions and greater quality of medical care. Longevity improvements in recent decades have been unprecedented, and affected both males and females, with males showing slightly greater longevity improvements than



females. Figure 15.1 illustrates changes in life expectancies for males and females in the US over the past four decades. It is expected that the figures based on the 2010 census (likely to be released in 2011) will show continuing improvement.

The general pattern has been common to most developed countries, even though the pace of longevity improvements has not been the same in all countries and not always steady. Changes in longevity and mortality have also differed, often greatly, from one subset of the general population to another. These differences represent potential dangers in pegging longevity improvements of pension plan participants or another population to that of the general population, and point to the need for carefully designed longevity indexes.

While longevity improvements can be seen as an overall trend, it is unclear whether and to what degree this trend will continue. The significant uncertainty associated with the magnitude of future longevity improvements is the primary source of longevity risk. General medical advances and breakthroughs in the treatment of cancer and heart disease could lead to a gradual acceleration or sudden jump in longevity improvements. On the other hand, the pace of longevity improvements could slow down instead of accelerating. (There is a minority opinion that the obesity epidemic in the US might even lead to decreases in life expectancies despite improvements in medical care.) Longevity risk is very difficult to quantify, especially since pension plans and most other holders of this risk are concerned with long time horizons, making any projections significantly more difficult.

Modelling longevity improvements

A number of approaches used to model future mortality incorporate longevity improvements. Chapter 13 describes some of these mortality models, including the popular Lee–Carter model and some improvements to it, as well as the application of the Markov process based on Brownian motion with a non-constant drift.

The P-spline (or penalised splines) model has become relatively popular in recent years. It too allows simulation of future mortality rates (through the Poisson process simulation of the number of deaths). The Cairns, Blake and Down (CBD) model, and its modifications suggested by several longevity researchers, uses an implicit assumption of a functional relationship of mortality rates across ages (Cairns *et al* 2006, Cairns *et al* 2007), differentiating it from most traditional models. Certain advantages of the generalised Smith–Olivier mortality model might allow its wider future use in stochastic

modelling of long-term mortality trends. The use of expert opinion and epidemiological models that reflect the effect on mortality rates of various mortality causes can improve the accuracy of the analysis, or at the very least add another perspective to the view of future longevity improvements.

Models that are only “backward-looking” and base all projections on the historical data are fundamentally deficient in making the assumption that historical mortality rates and trends contain all the information on the future behaviour of mortality rates. The most promising approaches are those that have these models incorporate some degree of expert opinion on potential future developments. Incorporating data not based on historical observations has always made actuaries uncomfortable due to the unavoidable degree of subjectivity needed for choosing and utilising such inputs. The approaches incorporating these inputs have, however, a strong potential to provide a more accurate probabilistic picture of future mortality and longevity.

The models do not reduce the uncertainty but allow one to better quantify it. They are extremely sensitive to inputs and assumptions. The output of a model is stochastic rather than deterministic: a multitude of scenarios are generated for mortality and longevity, representing the range of possible outcomes.

It is noteworthy that, in the past, practically all longevity forecasts in the developed countries have consistently underestimated the actual (realised) longevity, further highlighting the degree of uncertainty and financial risk associated with longevity improvements.

NATURAL HEDGES

Insurance companies can have longevity hedges already in place by simultaneously writing life insurance and annuity contracts. Greater-than-anticipated longevity will result in annuity payments being made longer than expected. At the same time, greater longevity means lower mortality, resulting in fewer life insurance claims and greater profits from the life insurance book. In theory, the two could offset each other. In reality, this type of hedge would never be perfect due to differences in the mortality characteristics of the populations of insured and annuitants. Still, this natural hedge does take away some of the longevity risk of an annuity writer. At the same time, it decreases the risk of greater-than-expected mortality affecting the life insurance book, since the negative financial effect of greater mortality on the life insurance book is offset, at least to some degree, by the positive financial impact of lower longevity on the annuity book.

Careful assessment of the relationship between the mortality rates of the insured and annuitants can allow a company to estimate the effectiveness of this hedge; in some cases, it may be possible to structure a swap between a life insurance writer and an annuity provider to take advantage of this natural type of hedge between life insurance and annuity products.

PRIMARY MECHANISMS OF LONGEVITY RISK TRANSFER

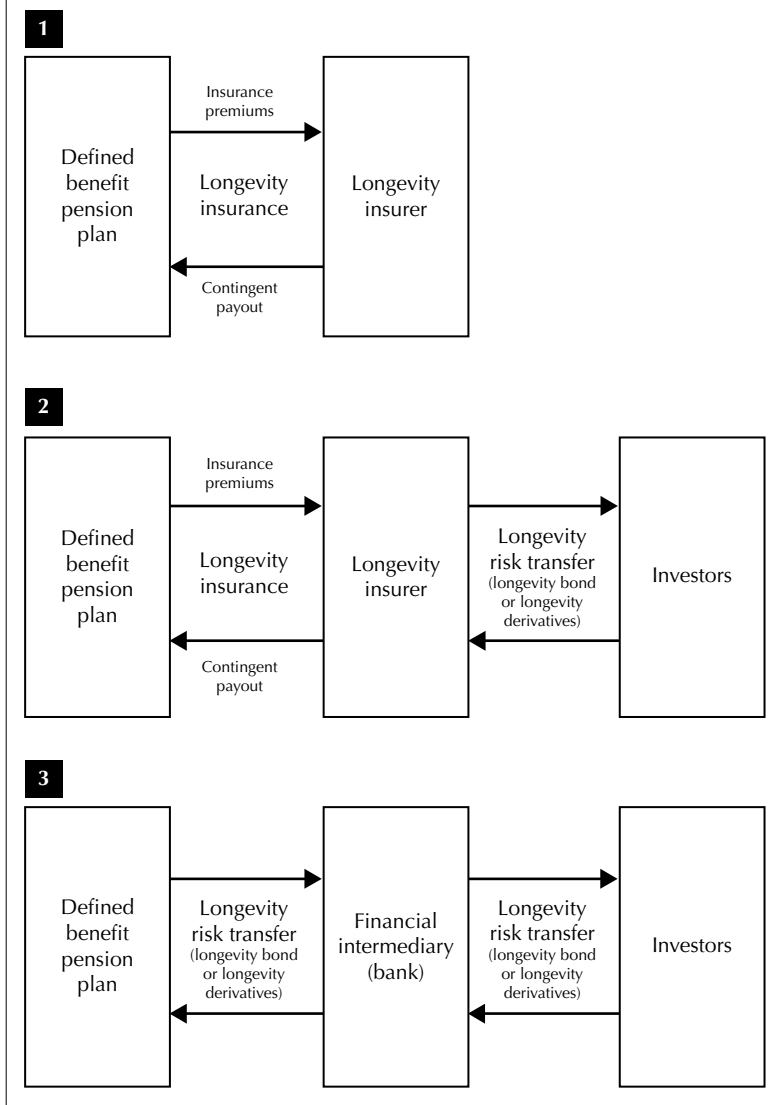
For annuity products, longevity risk used to be effectively handled by reinsurance. With the demand for longevity hedges growing, and limited reinsurance capacity, this solution can address the problem only to a very small degree, and it does not at all help to mitigate this risk for pension plans, where the need is greatest.

Transferring longevity risk to the capital markets presents a natural solution. Figure 15.2 shows three possible solutions to the problem of hedging longevity risk that are available to defined benefit pension plans. The first solution is obtaining an insurance policy covering the risk of longevity – that is, of plan participants living longer than assumed and the plan needing to make payments for a longer period of time. The insurance policy would cover any shortfall resulting from longevity being greater than the level specified in the policy. Pension Corporation in the UK was probably the first to introduce longevity insurance for DB pension funds, and it remains a very sophisticated player in the market. In the first solution shown on Figure 15.2, the insurance company retains the longevity risk that is being supported by the company surplus (shareholder equity) obtained through selling stock to investors and retained earnings.

The second solution is identical to the first one from the point of view of the pension plan. The difference is that the insurance company does not retain the risk but rather passes it along to the capital markets. The ways to effect this transfer are shown in Figure 15.2. Alternatively, the longevity insurer could retain some of the longevity risk while passing the rest of it to investors. This could be done because the company has sufficient capital to support some of the risk and there is no need to buy protection for that part of the risk from the capital markets. It could also be the case that the hedge the insurance company obtains from the capital markets is not perfect, and the insurer retains the resulting basis risk. The insurer may also choose to aggregate longevity risk from more than one source, achieving greater scale and diversification, before passing all or some of it along to the capital markets.

The third solution shown does not involve an insurance company but

Figure 15.2 Primary mechanisms for protection of defined benefit pension plans from longevity risk



rather has the pension plan enter directly into a financial contract to transfer longevity risk to the capital markets. In practice, a bank or another financial intermediary will enter into such a contract with the pension plan and assume the longevity risk. The bank would generally not retain the risk but rather would pass it along to the capital markets.


The financial instruments for transferring longevity risk to the capital markets are longevity derivatives and longevity bonds. Below we analyse each of these instruments.

LONGEVITY SWAPS

A longevity swap exchanges payments linked to predetermined fixed longevitys for those of the longevitys actually realised. Such a swap protects a DB pension plan from longevity risk by always providing the amounts that the plan needs to pay its participants in any given year, in exchange for the plan's making predetermined payments based on the expected longevity. If the actual and expected longevity are the same, the two sets of cashflows cancel each other out, and no actual cash exchange between the pension fund and the swap counterparty (the investor) ever takes place. On the other hand, if longevity turns out to be greater than assumed, resulting in the need to make higher-than-expected payments to the plan participants in a given time period, the shortfall will be made up by the swap provider.

Figure 15.3 provides an illustration of how such a swap might work. The payments shown on the top are calculated based on the expectations of future longevity of the pension plan participants as of the date the plan entered into the swap arrangement with a hedge provider. These payments are predetermined at the inception, and are made by the DB pension plan to the hedge provider. The payments shown at the bottom are the actual payments that the pension plan needs to make to its participants. They are not known in advance and are based on the actual, or realised, longevity of the plan participants, which differs from the longevity assumed when the parties entered into the swap agreement. The hedge provider makes these payments to the pension plan. The net payments made are shown in black.

In this illustrative example, the hedge proves its value, since the longevity in the later years ends up higher than expected, and the pension fund receives additional cashflows from the swap provider. The later years are the ones where the uncertainty is greatest and the longevity risk most significant. However, there is also a payment to the plan only a couple of years into the contract. While in this example the pension fund receives additional cashflows to cover the shortfall – based on the actual longevity being generally higher than assumed at inception – there is one point when the net payment is made from the fund to the swap provider, illustrating that the longevity can also be lower than assumed;⁴ in fact, all the payments might end up being made by the pension fund. This possibility by no means



detracts from the value of having the hedge in place, since its purpose is not to provide return to the pension fund but simply to protect the plan from longevity risk.

In a swap arrangement, both parties are likely to be required to post collateral, the amount of which might change over time according to a predetermined time schedule or the observed mortality levels in the future. The collateral requirement might be waived or the amount of required collateral reduced if the hedge provider is a highly rated counterparty.

The simplified example in Figure 15.3 has the swap based on the actual longevity of the pension plan members. It could be beneficial to have as a reference point the longevity of the general population rather than that of the pension plan participants. Such a derivative is easier to structure and has greater appeal to potential investors. If the market grows, there could be an opportunity to trade these derivatives, creating liquidity. The standardisation itself would likely contribute to the market growth. Standardised

solutions also tend to be cheaper. The hedge effectiveness does diminish when the general population is chosen as a reference point rather than the pool of the pension plan participants. It could be argued, however, that the difference is not particularly significant. While the longevity of the general population can vary significantly from that of the members of a specific pension plan, the changes in longevity – in particular, longevity improvements that represent the risk – are closely correlated, since they are all based on the same underlying trends. Dependence of hedge effectiveness on the choice of the longevity index versus the actual longevity of the plan participants needs to be carefully examined in each case, and benefits and disadvantages properly assessed.

It might be in the best interests of a pension plan to have a hedge only against significant – above a certain level – deviations of actual cashflows from expected amounts, due to longevity improvements. Such a solution is probably cheaper; it is analogous to having a deductible structure when purchasing longevity insurance.

MORTALITY FORWARDS AND SURVIVOR FORWARDS

While the type of longevity swap illustrated in Figure 15.3 can provide an effective cashflow hedge to a pension plan to protect it against future longevity improvements, it is a complicated bespoke instrument that is unlikely to be traded in the capital markets. To make longevity trading possible, simpler instruments should be created. “Simpler” does not imply a drop in hedge effectiveness, since the instruments can be used as building blocks to construct sophisticated hedges that mimic the longevity behaviour of a reference population with a significant degree of precision. It might even be possible to decompose the longevity swap illustrated in Figure 15.3 into some smaller and separately traded building blocks.

Some such building blocks are examined below, in particular q-forwards and survivor forwards that have been developed by the LifeMetrics team at JP Morgan and the Pension Institute.⁵ (The focus of LifeMetrics has been on mortality forwards rather than survivor forwards, but the two are related.)

Survivor forwards


A survivor forward is a derivative contract linked to a survival rate at a certain point t in the future. It is often referred to as “s-forward” in the terminology used by LifeMetrics and the Pension Institute. The contract itself is a swap, with only one payment made at its maturity. Effectively, a fixed survival rate is being swapped for the actual survival rate at maturity. The

payment under such a contract, performed at point t , depends on the difference between the actual survival rate of the reference population and the survival rate predetermined at inception. If the actual survival rate ends up being higher than the fixed rate based on initial expectations, the hedger receives a payment proportional to the difference. This payment then fills all or some of the funding shortfall resulting from greater-than-anticipated longevity. Conversely, if the actual longevity rate at time t ends up being lower than was originally expected, the hedger makes a payment to the swap counterparty.

The survivor swap described earlier can be replicated by a series of survivor forwards. In other words, survivor forwards can be used as building blocks in constructing a longevity hedge. However, survivor forwards themselves are not the instruments that can be most easily traded in the market. Simpler instruments, such as mortality forwards described later, might be better candidates for such tradable securities.

The payout of a survivor forward is shown in Figure 15.4 as a function of the actual (realised) survival rates. If the actual survival rate at the end of the period equals the fixed survival rate agreed on at the beginning, no payment exchanges hands. Greater-than-expected (fixed) survival rate results in a payment to the hedger.

Figure 15.4 Net cashflows of a survivor forward



Source: LifeMetrics and the Pension Institute


The payment illustrated in Figure 15.4 is shown on a net basis. On a gross basis, the cashflows at maturity involve the exchange of a payment proportional to the fixed survival rate for a payment based on the realised survival rate. Figure 15.5 shows the schematics of the cashflow exchange at maturity for a survivor forward.

The hedger, which would likely be a pension plan or an insurance company writing annuity contracts, makes a payment proportional to the fixed survival rate agreed upon at the inception of the contract. This cashflow is being swapped for a payment from the counterparty providing the hedge, with the payment being proportional to the realised survival rate at maturity. The payments are calculated as the notional amount times the survival rate. If the actual (realised) survival rate is greater than the expected (fixed) rate, the hedger makes a smaller payment than the hedge provider, thus receiving, on a net basis, positive cashflows at maturity. In the case of a pension plan, this payment provides the plan with extra funds needed due to greater than expected longevity improvements.

To mitigate the counterparty credit risk, collateral requirements would usually be part of the survivor forward contract. The amount of the collateral may be fixed or may vary depending on the time to maturity and the divergence being experienced between the implied expected and the actual survival rates. For a highly rated counterparty, no collateral might need to be posted; the contract would then specify the amount of collateral needed to be posted for each downgrade.

The term “survivor forward” is used because of the clear analogy with commodity or foreign exchange forwards. If we try to draw an analogy with interest rates, the equivalence is with spot rates rather than forward rates. The fixed longevity rate can be interpreted as the spot rate at the inception of the contract. Later, it loses this meaning and is seen as simply a survival-rate level referenced in the survivor forward contract. However, even at a

Figure 15.5 Schematics of the cashflow exchange for a survivor forward



Source: LifeMetrics and the Pension Institute

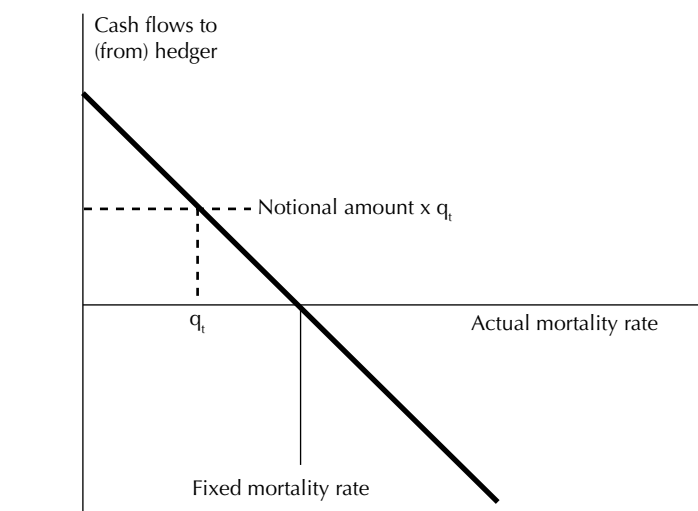
later point in time it can be translated into a spot rate from that point in time, to time t , based on the realised mortality up to that point. The analogy with commodity or foreign exchange forwards is more appropriate for survivor swaps.

Mortality forwards


A mortality forward (q-forward in the terminology of JP Morgan's LifeMetrics) is a swap contract exchanging expected or otherwise predetermined (fixed) mortality for the actual (realised) mortality. The payment is made only at one point in time – at maturity. In the case of a mortality forward, the cashflows are in the directions opposite to those for a survivor forward. The fixed payment, proportional to the fixed mortality rate, is made by the hedge provider to the hedger. The hedger is the one who makes the contingent payment, the value of which is proportional to the realised mortality rate at maturity.

On a net basis, if the realised mortality rate is lower than the expected one, a payment is made to the hedger. In the example of a pension plan, lower mortality means greater longevity and the need to make additional pension payments. The positive cashflow from the hedge would provide extra funds to make these payments. Greater-than-anticipated mortality will result in

Figure 15.6 Net cashflows of a mortality forward



Source: LifeMetrics and the Pension Institute

Figure 15.7 Schematics of the cashflow exchange for a mortality forward

Source: LifeMetrics and the Pension Institute

the positive net cashflow being received by the hedge provider. In this situation, the longevity is lower than expected, and the pension plan has lower-than-anticipated payments to make to its members.

The mortality rate in mortality forwards is not the cumulative rate from the inception of the contract to maturity. Instead, it is the regular annual mortality rate at maturity. In the LifeMetrics framework, this rate could be based on the index data provided by LifeMetrics. Within the same framework, however, it could be any other mortality rate chosen by the two swap counterparties.

Term structure of longevity (mortality) rates

Following the analogy with interest rates, a set of spot mortality rates (that can be determined based on the market pricing of survivor forwards) can allow us to calculate forward rates for each year in the period, which are simply the expected annual mortality rates in the future. This calculation is equivalent to the one that can be performed to determine expected future one-year spot interest rates, which are equal to the forward rates for those years.

The calculation works for interest rates under the expectations hypothesis. It does not work under the liquidity preference hypothesis. While in the absence of an active market there are no observable spot rates for mortality or longevity, it is likely that the liquidity preference hypothesis is applicable to the mortality term structure as well. In other words, it is possible and probably likely that a spot mortality rate is not equal to the product of the expected annual mortality rates over the time period. This suggests caution in trying to find arbitrage opportunities by comparing mortality rates implied by the market prices for survival forwards and mortality forwards. Such issues will need to be examined if there develops an active market in longevity. It is likely that the future forward mortality rates differ from the expected future spot

mortality rates by the amount of mortality risk premium effectively paid by the hedger to obtain protection against longevity risk.

Standardised index-based longevity hedges

JP Morgan, a recognised leader in longevity risk transfer, has developed its LifeMetrics indexes and methodology to facilitate the development of the liquid longevity trading market. LifeMetrics indexes can become industry standard, and the q-forwards (mortality forwards) based on the LifeMetrics methodology have a potential to become the tradable building blocks for hedge construction.

The analysis based on LifeMetrics tools has shown that, to obtain an effective longevity hedge, we need only a relatively small number of these building blocks, and that q-forwards can combine a range of ages (such as 40–49, 50–59, 60–69, 70–79 and 80–89) for each gender and still be an effective hedge when used in proper combination with weights appropriately chosen. In fact, LifeMetrics makes this process relatively easy and allows us to determine the best longevity hedge for pension liabilities based on these building blocks. JP Morgan believes that a very small choice of q-forward maturities is required. Taken together, all of the above translates into a relatively small number of the q-forward contracts that would be needed, making it easier to establish a market for trading these instruments.

The determination of hedge effectiveness is critical in the choice of the best longevity hedging mechanism for a pension plan. Concepts such as longevity value-at-risk (longevity VaR) have been used to describe the longevity risk of a pension plan before and after applying a longevity hedge. The goal is not necessarily to have a perfect hedge but rather to have the most cost-efficient solution that reduces longevity risk to an acceptable level. Standardised hedges are cheaper and easier to implement; at the same time, if properly constructed, they can have a rather high degree of hedge effectiveness. So far most of the longevity risk transfer solutions included customised rather than standardised index-based hedges. As the marketplace becomes more comfortable with the new tools and as the ability to properly quantify the risk and determine hedge effectiveness improves, it is likely that the standardised hedge solutions will become more popular.

LONGEVITY BONDS

The idea of transferring the risk of longevity to the capital markets by means of a longevity bond is not new. This appears to be a natural way to transfer longevity risk, but the implementation is not easy. In this structure, the bond

provides investors with declining cashflows linked to longevity of the population whose longevity risk is being hedged. The link might not be direct, and a proxy for longevity of a specific population may be used.

The following main types of longevity bond structures have been proposed.

- ❑ Zero-coupon longevity bonds that make only one payment to investors at the end of the term, with the amount of payment being linked to a longevity index of a population. A longevity risk hedge would involve a combination of such bonds.
- ❑ Fixed-term, or regular, longevity bonds with coupon payments being tied to longevity experience of a population, with coupons generally declining (at least after a certain point), due to the mortality of the reference population. The coupons might initially increase by design, if the pension plan payments are expected to increase in subsequent years due to additional participants reaching retirement age. Another reason for potential temporary increases in coupons might be fluctuations in longevity experience of the chosen index, in particular when it reflects the actual longevity experience of participants in a small pension plan.
- ❑ Open-term longevity bonds that are different from the fixed-term bonds in that the coupons are paid as long as there are individuals alive in the reference population. The maturity of such a security would not be known in advance and is a stochastic variable. In practice, for such a bond to ever be placed – which is probably unlikely under any circumstances – there should be a mechanism for limiting the term by, for example, making a bigger last payment if the survivor index falls below a certain level. The maturity would still be unknown in advance even in this case. Such a bond has been referred to as a survivor bond, but the terminology is inconsistent since the same term has been applied to the traditional, fixed-maturity longevity bonds.
- ❑ Principal-at-risk longevity bonds with coupons fixed at issue – but not necessarily level along the term of the bond – while the amount of the last payment (principal) is linked to a longevity index. Such a bond provides more of a value hedge rather than a cashflow hedge against the risk of increased mortality.
- ❑ Inverse longevity bonds are the opposite of regular longevity bonds in that they have coupon amounts rising rather than falling over the term, with an inverse relationship between the coupon amounts and the value of a longevity index. These are actually mortality rather than longevity

bonds; they are mentioned here because they have a potential to become part of the longevity risk transfer toolkit when used in combination with regular longevity bonds. Under certain conditions such a combination can replicate the cashflows of a traditional (not linked to longevity or mortality) bond.

Using the building blocks outlined above, or the general ideas used in their construction, we can devise a number of other longevity-linked bond-type structures. The concept of balloon maturity can be easily applied to these structures. CDO-type structures (collateralised longevity obligations) can be created if the market ever becomes sufficiently large. Choosing the optimal structure is very important because so far longevity risk transfer in the form of longevity bonds has not been successful, even though there are now renewed efforts to structure such bonds.


BNP Paribas, in 2004, structured the first longevity bond on behalf of the European Investment Bank (EIB). This bond serves as an important reference point for structurers of longevity risk transfer instruments. Ultimately, the bond was a failure in the sense that BNP Paribas was able to find investor interest for only a small part of the proposed issue. Despite the failure – and in part because of it – this attempt represents an important stepping stone in the development of the longevity risk transfer markets, and valuable lessons can be learned from it. These lessons are valuable because, even though the recent activity in longevity risk transfer has been focused on derivative instruments, the appearance of longevity bonds using improved structures appears to be likely.

The BNP Paribas EIB longevity bond

The bond was supposed to have the total value of approximately £540 million, or €775 million, and the tenor of 25 years. Investors in the bond were expected to be pension funds, and the cashflow structure was intended to approximate the effects of changes in longevity on DB pension plan payments. While the European Investment Bank was the issuer, BNP's role was that of the structurer, marketer, manager and book-runner.

In this structure BNP assumes longevity risk from EIB, and later reinsures it to Partner Re. There is also an agreement between EIB and BNP to swap sterling and euro payments.⁶ The notes are not officially rated, but effectively they receive the rating of the issuer, which is AAA. The EIB has credit risk exposure to BNP, and BNP in turn is exposed to the credit risk of Partner Re. The credit risk in the structure is important to the parties that

Figure 15.8 Cashflow for a longevity bond: the example of the BNP Paribas/European Investment Bank structure



might have to bear it, but it is of no relevance to investors since they are not exposed to this risk.

Cashflows between the bond issuer and the investors are reflected in the structure shown in Figure 15.8. Investors, which in this case are pension funds, purchase the bond and thus provide cashflow to the issuer at time $t=0$. The coupons, paid annually for 25 years, decline based on a chosen longevity index. There is no principal repayment at maturity, and there are no embedded options.

The bond payout at the end of year t from the issue equals £50 million times the cumulative survival rate in the initial cohort at time t . The cumulative survival rate, CSR_t , was defined as the proportion of survivors at time t for the cohort of males aged 65 at issue⁷ based on the English and Welsh general population mortality data as reported by the government.⁸ Using the terms defined in Chapter 13, the cumulative survival rate can be calculated as


$$CSR_t = {}_tP_{65}^{male}, \text{ or } CSR_t = \frac{S^{male}(65+t)}{S^{male}(65)}$$

where the probability of staying alive at the end of year t from issue and the survival function are based on the cohort of 65-year-old males at issue. The actual payments would of course differ from the expected. The cumulative survival rate can also be calculated as

$$CSR_t = \prod_{i=1}^t (1 - q_{65+i-1}^{male})$$

where again the mortality rate is that of the cohort of English and Welsh males who were 65 years old at issue. CSR_t is a random variable; it is observable only at time t from the issuance.

Figure 15.9 shows the projected cashflows for the bond based on the government projections of mortality rates at the time BNP Paribas was marketing the bond to investors.

Figure 15.9 Projected annual bond coupon payments

In this structure, greater-than-expected cohort longevity results in higher coupon payments, while lower-than-expected longevity leads to smaller payments.

So far, due to longevity improvements the realised longevity for this cohort has been greater than was projected at the time. Higher longevity would mean coupon payments greater than shown in Figure 15.9 and likely losses for Partner Re, the company that agreed to reinsure the longevity risk assumed from EIB by BNP Paribas.

Lessons from the failure of the BNP Paribas / EIB longevity bond

The bond failed since it did not generate sufficient demand. What were the reasons for this, and do they have to do with the structural issues that could be addressed? Or are there some fundamental flaws that make longevity bonds in general a wrong instrument for transferring longevity risk to the capital markets? The main reasons for the failure were the following.

- ❑ Pension funds, who were the target investors, did not perceive longevity as a significant risk and did not believe it was cost-effective to hedge. This reason for failure was not specific to longevity bonds but probably would have applied to any instrument for longevity risk transfer. This view has been changing and there is now greater awareness of longevity risk and its potential implications. Continuing the educational process and, even more important, better ways to quantify the risk will likely overcome this difficulty.
- ❑ No regulatory benefits would have resulted from hedging longevity risk. This reason too was not specific to longevity bonds but would have applied to any instrument for longevity risk transfer. Pension funds in the

UK use actuarial tables that are updated infrequently and do not include future mortality improvements. The views of regulators have been changing and the attention paid to the longevity exposure of pension funds growing. The new regulatory regimes are expected to require, at least at some point in the future, proper management of all risks, including the risk of longevity, in many jurisdictions. Unfortunately, at this point most jurisdictions, including the US, pay virtually no attention to the risk of longevity in DB pension plans.

- ❑ There existed real or perceived issue of basis risk arising from the differences between the actual longevity experience of the participants in a particular pension plan and that of the index based only on males aged 65 at issue in the general population of England and Wales. This reason too is not unique to longevity bonds and could equally apply to longevity derivatives. The concern was justified, especially because quantification of the basis risk was difficult. Since then, however, better modelling tools have been developed, making it easier to assess hedge effectiveness and make informed decisions. Development of advanced tools for stochastic modelling of mortality and longevity, including longevity improvements, is continuing. Two ways to address the issue of basis risk are: (1) the use of actual longevity experience of pension plan participants instead of any other index; and (2) the use of a more sophisticated index that would more closely mimic the composition of the population of the pension plan participants. The latter might be preferable for market growth since it allows a degree of standardisation that can facilitate trading. However, basis risk will always be greater in this solution than in the first one.
- ❑ At the time, pricing tools had not been sufficiently developed. There existed a concern on the part of investors that the bond was overpriced. In fact, the application of the improved modelling tools seem to show the opposite, that the bond was underpriced. The actual experience certainly shows greater longevity improvements than were assumed in pricing. In retrospect, the concern was not justified, and now there are better tools and methodologies for assessing and pricing the risk. No changes to the bond structure are necessary to address this concern.
- ❑ Hedging longevity risk by purchasing a longevity bond requires a relatively significant upfront capital expenditure. Longevity derivatives would typically not require that level of an upfront expense.
- ❑ While not a reason for failure of this particular bond, the structure did not contain direct transfer of the longevity risk to the capital markets. Rather, the ultimate risk bearer was a reinsurance company. Despite the possible

argument that reinsurance companies are themselves supported by capital markets, this limitation could be seen as a potential flaw in the structure of the bond. Several ways to redesign the bond are available. One of them is to implement a structure where the longevity risk hedger is an issuer (or sponsor if a special purpose vehicle issues the bond) rather than a holder of a longevity-linked security. There are also ways to address this concern in the existing structure. In general, the fact that a reinsurance company is ultimately providing longevity risk protection should not prevent longevity bonds from being issued. It is a flaw only in the sense that it could make the market growth difficult beyond a certain level, once the reinsurance industry capacity has been exhausted.

It appears that longevity bonds have no unsolvable structural issues. However, it remains a fact that, while longevity derivative transactions have been performed, no longevity bonds have yet been issued. This may mean that, even though longevity bonds present a solution to the transfer of longevity risk, other solutions are simply more efficient. This may very well be the case, but it is still likely that longevity bonds will be issued in the future; perhaps they will become the preferred longevity risk transfer instruments in some situations while derivative instruments will be more appropriate in others. The fact that investment bankers are now having active discussions with clients about issuing longevity bonds suggests that these securities will be used for longevity risk transfer, possibly very soon. Longevity insurance and reinsurance will certainly exist as well, but the insurance capacity is limited unless the longevity risk is then transferred to the capital markets, again in the form of either longevity derivatives or longevity bonds.⁹

Comments on longevity bond pricing

Pricing of a longevity bond is based on calculating the net present value of the expected cashflows in the probabilistic framework. The primary uncertainty – future survival rates – has to be modelled stochastically. This takes us back to the question of modelling longevity improvements and the various approaches that can be used for that purpose. It is interesting that in addition to the numerous traditional methods for stochastic modelling of mortality rates, some have utilised the Wang transform for pricing longevity bonds and other longevity-linked securities. The Wang transform was briefly introduced in Chapter 3 as one of the ways to price property cat bonds. It utilises the distortion operator to arrive at the “distorted” cumulative density function $F^*(x) = \Phi[\Phi^{-1}(F(x)) + \lambda]$, where Φ is the standard normal

cumulative distribution and λ is the market price of risk. In this case, $F(x)$ can be the mortality rate. The difficulty arises in determining the λ parameters. One of the approaches suggested to calculate λ 's is to derive them from the known market prices of annuities. This approach is based on the assumption that investors would agree to prices based on the same transformed distribution as in the annuity pricing. This assumption has not been validated.

The example of the previously described Wang transform as a pricing method for longevity-linked securities is brought up for the sole purpose of showing the wide range of pricing approaches. The more natural methods are based on stochastic mortality modelling; some of them have been mentioned in the section on longevity improvements above and in Chapter 13. It is worth noting that the freely available LifeMetrics tools, while not designed specifically for pricing longevity bonds, include software to allow the use of several stochastic mortality models.¹⁰

MORE ON OTHER SOLUTIONS FOR LONGEVITY RISK MANAGEMENT IN A DB PENSION PLAN

While the discussion has been focused on direct transfer of the longevity risk to the capital markets, other solutions exist as well. In many cases, they might be preferable to the use of longevity derivatives or longevity bonds. Some of these solutions are outlined below.

□ **LONGEVITY INSURANCE.** As mentioned above, longevity insurance is now available to transfer the risk of longevity from a pension plan to a longevity insurer. Pension Corporation was the first to introduce this product, but now a number of companies offer longevity insurance. This type of insurance protects pension plans against the risk of having to make payments to the plan participants due to their living longer than expected. It reimburses the pension plans for the extra cost associated with the additional payments to the pensioners. A longevity policy is highly flexible and can reflect the exact population – member by member – of the pension plan participants. Longevity insurance provides protection only against the risk of longevity, but this risk might be the biggest for many pension plans. Side-by-side solutions, including protection against other risks, can also be implemented, using the same longevity insurer or another party to provide the hedge. For example, an inflation hedge can also be provided – sometimes in the form of an insurance policy, but more often as an inflation swap.

- ❑ **BUY-IN CONTRACTS.** Buy-in is a comprehensive insurance solution involving taking over all liabilities of a pension plan. It allows the de-risking of pension plan liabilities, while the investments continue to be held by the pension trust, and the pension plan continues all administrative functions. The risks transferred include longevity, inflation, interest-rate and other investment risks. Buy-in is similar to the plan buying a bulk annuity from an insurance company. Both assets and liabilities remain on the pension plan balance sheet, but the liabilities are hedged. As in longevity insurance, the pension plan is exposed to the counterparty risk of the insurance company. This has led to the demand, in some cases, for mitigating credit risk by segregating (ring-fencing) the assets received from the pension plan (buy-in price) and often holding them in a separate trust account as a collateral.
- ❑ **BUY-OUT CONTRACTS.** Buy-out is an even more comprehensive solution than buy-in. In this case, the insurance company takes over both assets and liabilities. The plan sponsor and the trustees relinquish all their responsibilities, which are in turn assumed by the insurance company. The plan sponsor (employer) no longer has any responsibilities with regard to payments under the pension plan, and any related liabilities are removed from its balance sheet. (In some cases, the buy-out is done only for some classes of the pension plan participants. The plan sponsor and trustees then retain their responsibilities for the other classes of pension plan members.) The administration of benefits is no longer the responsibility of the pension plan trustees but is done by the insurance company or its agent. Buy-out cannot include future benefit accruals.

There could also be partial solutions such as partial buyout. Pension Corporation in the UK offers a so-called pension plan sponsorship. This solution makes Pension Corporation the owner of the pension fund, but the backing of the original plan sponsor is not removed. The assets of the pension plan also remain in place. There is no insurance contract in the beginning, and Pension Corporation does not generate any returns unless and until the pension benefits of the plan participants are protected by an insurance policy. Solutions incorporating elements of a partial buy-out and the traditional liability-driven investing (LDI) also exist.

Insurance solutions may best address the needs of pension plans to hedge the risk of longevity as well as other risks. However, using insurance instead of directly accessing capital markets does not take away from the need for capital markets solutions to the problem of longevity risk. Insurance and

reinsurance companies may aggregate longevity risk, but, as the amount of risk grows, so does the need to transfer it to the capital markets. Doing so by way of issuing additional equity is not the most efficient solution. Statutory requirements concerning minimal capital levels, as well as the need to maintain certain ratings, make it less capital efficient to keep all the risk on the insurer's balance sheet, compared with transferring it to the capital markets in a more direct way by issuing longevity-linked securities.

INDEXES OF LONGEVITY

Transparent, reliable longevity benchmarks can promote market growth; they are a prerequisite for the creation of a liquid market in longevity risk transfer. Having a choice of longevity indexes can minimise basis risk and increase hedge effectiveness for those seeking to offload longevity risk. Those wishing to invest in and trade longevity risk also require a reliable reference point and the degree of standardisation that can come only with properly constructed longevity indexes.

Creating an index of longevity and mortality is an important but difficult task. Basis risk and data reliability are just two of the issues to consider. Many believe that, if a standard index – a measuring yardstick accepted by the whole market – existed, it would contribute to the ability to create and trade in instruments of mortality or longevity risk transfer.

Despite the obvious basis risk issues, the introduction of a standardised measure for mortality and longevity is useful if we are ever to see a liquid, tradable market in mortality and longevity risk, as opposed to private capital markets transactions. Having a transparent standardised measure of mortality risk enables the creation of mortality/longevity swaps, structured notes and other instruments. It also facilitates the settlement of such contracts.

Credit Suisse Longevity Index

In an attempt to create such a standard measure, in 2005 Credit Suisse introduced a simple index designed specifically to facilitate structuring of capital markets instruments for the transfer of mortality and longevity risk. Called the Credit Suisse Longevity Index, the index is based on US government data collected by the Centers for Disease Control and Prevention. Any actual portfolio would present a composition by age and gender different from the general population mix, adding to the basis risk in any mortality or longevity risk transfer transaction based on this index. This additional risk could be decreased by using a combination of sub-indexes included in the

Credit Suisse Longevity Index. The index includes sub-indexes for attained ages of 50, 55, 60, 65, 70, 75 and 80 separately for males and females, as well as a composite of both genders. This information allows us to construct a custom index by weight-averaging the data for different age-gender combinations, to more accurately mimic the actual insurance portfolio and to decrease basis risk. The index data also includes 30-year projections, based on the assumption that the age and gender population mix will remain constant during the entire projection period. It is also assumed that mortality improvements will continue at the historical rate.

The index is no longer published by Credit Suisse and is not available to the general markets for use in structuring and trading longevity transfer instruments. It appears that, at least for now, Credit Suisse has stopped updating the index. Instead, other indexes have been developed and introduced to the marketplace. A discussion of these follows.

LifeMetrics Index

In 2007, JP Morgan launched its own index called LifeMetrics. The index covers four countries: the UK (limited to England and Wales), the US, the Netherlands and Germany. It is likely that it will be expanded to other countries as well if it becomes more widely used. The data underlying the index is collected by the government agencies, is independent and not subject to manipulation and is based on the broadest datasets available. The methodology is fully transparent and available to the public. The LifeMetrics indexes are part of the LifeMetrics Longevity Toolkit developed by JP Morgan. The toolkit, made available to the public, also includes software tools that can be used for developing mortality and longevity projections. LifeMetrics Longevity Toolkit was created by JP Morgan based on research assistance provided by leading researchers, in particular the Pension Institute. It includes tools for stochastic modelling of mortality and for making longevity projections. It also has tools for addressing the issue of basis risk arising from the differences between the longevity experience of the actual population of pension plan participants and that of the general population reflected in the LifeMetrics index.

Watson Wyatt serves as the calculation agent for the index. As is the case with any index based on government data, there is a lag in reporting. This lag depends on the country and is unavoidable. Indexes include crude and graduated mortality rates as well as period life expectancy.

To encourage the adoption of the index and its general methodology as an industry standard, JP Morgan has even offered to donate all rights to the

LifeMetrics family of longevity indexes to the Life and Longevity Markets Association (LLMA).

The LifeMetrics methodology developed by JP Morgan is probably even more valuable than the index itself. It gives market participants access to some of the best tools available for quantifying longevity risk, building probabilistic forecasts, and ultimately facilitating the growth of a liquid market in longevity. LifeMetrics represents a significant advance in the development of the framework, data and tools needed in the longevity risk transfer market.

Deutsche Börse Xpect Index

In 2008, Deutsche Börse introduced its own family of indexes. These now cover Germany, the Netherlands, and the UK (limited to England and Wales), providing age indexes and cohort indexes. The first one (Xpect Age indexes) is based on “open” populations by country and represents average life expectancies of defined age ranges. The calculation is based on a weighted average of all birth years within an age range. The age indexes are reported on an annual basis. Separate values are available for males and females. The second one (Xpect Cohort indexes) is based on “closed” populations by country and represents life expectancies of these cohorts, each of which includes a range of ages. The indexes are reported on a monthly basis.

Deutsche Börse can also design custom indexes (Xpect Portfolio indexes) that mimic longevity and mortality characteristics of existing portfolios of longevity and mortality risk, and are based on the other two indexes (Xpect Age indexes and Xpect Cohort indexes).

Xpect Data is a companion product and the source of data for calculating the Xpect Age, Xpect Cohort and Xpect Portfolio indexes. Xpect Data includes generational life tables that include mortality rates and life expectancies. The methodology for calculating mortality rates and life expectancies is disclosed.

Information is provided on a monthly rather than annual basis due to the incorporation in the data of some elements that do not come from central governments. Data from the central governments are updated very infrequently, and Deutsche Börse supplements this information with the more current data obtained directly from municipalities and other sources. While generally relatively transparent, this process includes a number of subjective factors and makes the data less easily auditable. An investment bank or any other entity with a potential financial stake in the longevity and mortality market would not be able to build a tradable index using data that is either proprietary (at least to some degree) or obtained from sources where there

could be an informational advantage for some parties. Deutsche Börse is able to avoid an appearance of the conflict of interest because it does not have a stake in specific longevity risk transfer deals but would only like to see their number grow and, ideally, to see active exchange trading of longevity instruments.

Other indexes

Other indexes have been proposed but have not received traction in the market place. It may be argued that some of them could be used for bespoke transactions. The attempt by Goldman Sachs to introduce QxX as a standard index for use in life settlements investing and hedging has not been successful. It is beneficial for the growth of the market to have fewer competing indexes so that industry standards can be established. It is also beneficial to have a smaller number of longevity risk transfer instruments to promote market liquidity. Of course, these instruments should be flexible enough to effectively manage the issue of basis risk and to improve longevity hedge effectiveness. The standardisation would not mean an elimination of bespoke solutions, but ideally these solutions will be based on the simple and separately tradable building blocks such as those developed by LifeMetrics.

INVESTORS IN LONGEVITY

While the identification of the main holders of longevity risk is relatively easy, with DB pension plans and life annuity providers being the obvious choices, it is less straightforward to identify the types of investors for whom getting paid for assuming longevity risk is most beneficial.

DB pension funds are not the best investors in this asset class since they are already long longevity risk, and adding longevity exposure by investing in longevity-linked ILS only increases this risk. Hedge funds are a natural choice, but only if investors in these hedge funds do not include pension funds. Allocating assets to alternative investments could be an important part of a pension fund investment strategy; there is a need to be careful, however, not to increase the pension plan longevity risk accidentally through an investment allocation to a fund that is long longevity risk.¹¹ On the other hand, a small allocation to longevity-linked securities might not have a material effect on the risk. This issue also has to be addressed in the hedge fund disclosure to investors. Currently, the problem is largely hypothetical since very few longevity-linked securities exist; as the market develops, the issue will grow in significance. Dedicated ILS funds fall into

the same category. Some of these funds may have the specialised expertise needed for the analysis of longevity risks.

The most natural longevity hedge providers are insurance companies that write life insurance – in particular, longer-term products such as whole-life and guaranteed level premium term life insurance. These companies are exposed to mortality risk and are short longevity risk, and would benefit from longevity exposure. The hedge effectiveness is not necessarily high, due to the significant differences in the mortality characteristics of life insurance policyholders and pension plan participants and annuitants; but investing in longevity will still reduce the mortality risk of these companies, in addition to providing possibly attractive investment returns. Insurance companies also have the advantage of actuarial staff and expertise in mortality analysis. Care should be taken, however, in trying to apply traditional actuarial methods and statistics based on life insurance mortality to the analysis of longevity and the probabilistic projections of long-term longevity improvements. In-house expertise might not be adequate to this task and could lead to a false confidence in being able to understand and properly model longevity risks. In general, the life insurance industry does not have much capacity for taking on the longevity risk of DB pension plans. The mortality (life insurance) and longevity (annuities) risk in the insurance industry are almost evenly balanced, with mortality risk being only slightly greater than the longevity risk for the industry as a whole.

Family offices are in a position to assume some longevity risk. Longevity-linked investments are not appropriate for most individual investors since they are exposed to the risk of their own longevity – that is, if they live longer than they expect, they face a greater chance of depleting their personal savings. Wealthy individuals are less subject to the risk of their savings being depleted, which is why family offices can take on longevity risk and profit from it. It should be noted, however, that family offices generally do not have the resources to develop expertise in longevity analysis. Should they decide to invest in longevity, the best way to do it would be through a specialist fund.

There are sectors of the economy – from pharmaceutical companies to nursing-home facilities – that can benefit financially from longevity improvements. However, they can rarely invest in longevity risk and would not consider it to be a hedge to protect the future profitability of their businesses.

Endowments do not seem to have any reason to avoid investing in longevity. With some exceptions, they are generally not long longevity risk

and so would benefit from investing in an asset class that likely provides exposure to exotic beta and a potential for relatively high returns.

While a number of potential investors have been identified in the above discussion, the market has not yet developed; and, while there has been some investor interest, the number of actual transactions has been small, undertaken by those who are best able to analyse the risk. There are, however, categories of investors for whom assuming longevity risk – given proper compensation – makes sense. In this light the longevity market has a strong potential to further develop and grow.

MARKET DEVELOPMENTS

Until recently, the longevity risk transfer market had seen a lot of general activity but very few actual transactions. This situation seems to be changing rapidly with the growth of longevity risk transfer deals in the UK. The UK is the first because of the changes in the regulatory environment and the fact that its DB pension liabilities are by far the greatest of all European countries. If and when the US follows suit, the size of the longevity risk transfer market could skyrocket; but the “when” may not come any time soon.

Most of the recent developments had to do with the pension plan buy-ins or buy-outs done mostly by insurance and reinsurance companies; and these include the transfer of other risks in addition to longevity. Pension Corporation, focused exclusively on this market, has been active in pension buy-outs and buy-ins but was also the first to develop a longevity insurance product for pension plans. Now several other companies are offering this product.

The development of the insurance part of the market rather than the direct transfer of longevity risk to the capital markets addresses the interests of hedgers by eliminating basis risk but does little to promote a liquid market in longevity. Ultimately, longevity insurers are likely to pass most of the aggregated longevity risk to the capital markets; but this has not happened yet.

Meanwhile, however, in the UK direct transfer of longevity risk to the capital markets has started to develop. Longevity swaps have been placed in the market, though almost all of them were based not on a standard index but, to eliminate basis risk, on the actual exposure of the pension plans. In 2008 Lucida plc, a specialised longevity insurer, hedged some of its longevity risk through a longevity derivative contract with JP Morgan linked to the LifeMetrics longevity index for England and Wales. It was not an insurance contract but a q-forward derivative, with ISDA and CSA documentation used. The transaction was fully collateralised.

Later the same year, JP Morgan assumed longevity exposure of £500 million in the UK pension liabilities of Canada Life; simultaneously, JP Morgan entered into longevity swap agreements with several investors to pass the longevity risk to the capital markets. The hedge in this case was not based on a standard index but rather on the actual longevity exposure of the pension plan. The 40-year cashflow hedge protected the closed portfolio of pensions from the risk of longevity improvements as well as any shortfalls due to random fluctuations in longevity rates.

Another very large transaction in 2010, the assumption by Deutsche Bank of the longevity risk of £3 billion in the UK pension liabilities of BMW, through its insurance subsidiary Abbey Life, has again demonstrated the market potential. It appears, however, that most of the risk has been passed along not directly to the capital markets but rather to the reinsurance companies.

These are just some of the examples of the actual transactions involving longevity risk transfer. A number of transactions have been done whereby a bank, through a subsidiary insurance company, provided longevity protection in the form of insurance, and then passed on the risk to the capital markets in the form of longevity derivatives.

Since the insurance and reinsurance companies are likely to reach their longevity risk capacity as the market continues to grow, the importance of longevity swaps and other instruments for direct transfer of longevity risk to investors is likely to increase. In 2010, a consortium of investment banks and insurance/reinsurance companies was formed to help facilitate longevity risk transfer and to develop standardised indexes for trading longevity and mortality risk. The consortium, called the Life and Longevity Markets Association (LLMA), is focused entirely on longevity risk transfer related to pension funds and to annuity providers rather than to life settlements products. If the LLMA is successful in the development of the relatively simple standardised products and reliable indexes that will gain general acceptance in the industry, the longevity risk transfer market will likely experience a significant boost to growth.¹²

EXTENSION RISK IN TRADED POLICIES

Life settlements investors, in their exposure to significant longevity risk, stand apart from the pension funds and annuity providers. Managers of life settlements portfolios have unique issues in hedging their risk of longevity being greater than expected. The populations of insureds who have chosen to settle their life insurance policies differ quite significantly from both the

general population and even the insured population with seemingly the same characteristics. They also differ from portfolio to portfolio and, more importantly, from one insurance policy origination source to another.

The risk of longevity improvements extending life expectancies is growing in life settlements, as there is a general growth of the proportion of policies with longer LEs. Such policies have a greater exposure to longevity improvements. In addition, traded policies have a disproportionate number of male versus female insured, and males at most ages and in almost all developed countries have been experiencing greater longevity improvements than females. This again makes life settlements more exposed to the risk.

However, the main longevity risk in life settlements is not that of unanticipated longevity improvements but of the LEs having been understated from the very beginning – the point when the policies were sold to investors – and of this underestimating still not being recognised in portfolio valuations. This risk is systematic and has to do with the way life expectancies have been (and to a significant degree still are) determined in the market. The process, described in Chapter 12 and touched on in Chapter 13 and Chapter 14, has led to widespread underestimating of life expectancies in life settlements; some of this underestimation has been corrected and some still has not. Since many portfolios of life settlements are small in size, at least in relative terms, fluctuations in performance are expected. When the actual death benefits for a portfolio are less than anticipated, it is often possible to discount the difference by attributing it to random statistical fluctuations around the mean rather than systematic underestimating of the mean itself. Valuation errors may persist for quite a number of years; currently, they are widespread in the marketplace.

Random fluctuations of realised longevity around the mean present a source of longevity risk that is of much greater magnitude in life settlements, due to the smaller sample sizes (number of insured lives in portfolios of life settlements), than in typical pension plans.

Managing longevity risk in life settlements

Currently, hedging options available to investment managers of life settlements portfolios are very limited and in most cases nonexistent. The first step to effective management is proper valuation of the policies – both at the time of purchase or sale and later, when the policies are part of the portfolio. The knowledge that systematic risk of underestimating longevity is present in life settlements should be a constant reminder, to portfolio managers and

analysts, of the need to revise and validate pricing assumptions based on both the performance of their portfolios and the information that becomes available from external sources. This systematic risk of the mean being understated is difficult to hedge effectively. However, given the inefficiency of the market, those with greater expertise might be able to sell and buy policies at prices that are not consistent with the degree of longevity risk of the policies, and to trim the number of the policies with a greater chance of the underestimated LEs while possibly even generating return from the trading. In mitigating the risk of systematic underestimating of life expectancies, a natural approach is to diversify the portfolio in terms of gender, age and medical condition of the insured lives; of types of policies; and, most importantly, of policy origination sources and LE providers involved. Assembling a bigger portfolio, even at the expense of moving to lower average face value of the policies, is another portfolio management tool that can reduce longevity risk associated with random fluctuations due to small portfolio size, as well as the longevity risk associated with possible overexposure to a “contaminated” policy origination source.¹³

Insurance and reinsurance have been used to transfer away the risk of longevity being greater than projected in portfolios of life settlements. Only a handful of such transactions have been performed. This type of longevity risk transfer is unlikely to grow and may completely disappear because on virtually every transaction the insurance companies have lost money.

Longevity derivatives in life settlements

Longevity derivatives tied to a general population index are of little use in life settlements. Life settlements longevity characteristics are too different from those of the general population for such a hedge to be effective. The correlation between longevity improvements of life settlements and those of the general population is relatively low.

An attempt was made by Goldman Sachs in 2007 to create a liquid market in longevity and mortality by introducing the QxX.LS (QxX) index, which directly references longevity experience of life settlements. This was done with the goal of facilitating trading in synthetic longevity securities. QxX, at 46,920 initially, is big enough to be representative of the general life settlements population, with the understanding that life settlements pools have a very significant degree of dispersion around the mean – within an individual pool and between the pools – that is greater than what is found in other segments of longevity risk transfer. The population in the QxX pool referenced the cohort of lives with individually identifiable information

stripped away. Transparency (in the calculation and in the choice to report the index on Bloomberg) and monthly tracking of the longevity performance of the pool were both intended so the index could facilitate derivative transactions in the life settlements market and enable the creation of synthetic life settlements securities. From the very beginning, the data and all calculations were publicly available and a third-party verification agent engaged. In addition to creating synthetic life settlements, the introduction of the index opened the door to constructing longevity derivatives referencing the pool performance; these were intended to become a way for investment managers of portfolios of life settlements to hedge some of their exposure to longevity risk, as well as to facilitate trading. The QxX index was well constructed also, in that it excluded a very controversial part of the life settlement market: that of insured lives suffering from HIV and AIDS. While lives with LEs less than two years were not excluded from the entering cohort, the HIV and AIDS exclusion still went a long way in distancing the index from the viatical market, with all the surrounding controversy. Such an index can be used not only for hedging mortality improvements but also for hedging systematic understating of LEs across the life settlement industry (possibly even realising an arbitrage opportunity, unless participants in the synthetic market already agree that the LEs are understated). In this case, the hedger will take a derivative position that benefits from the QxX reference pool longevity being greater than anticipated by the market.

The QxX index was intended to be only the first in a series of life settlement indexes introduced by Goldman Sachs. A year later, the QxX.LS.2 index was introduced to track the longevity performance of life settlements for individuals over the age of 65 with specific impairments that included cancer, cardiovascular conditions and diabetes. The initial size of the reference pool was 65,655. This index provided life settlement investors with the means to hedge the risk of longevity extension due to medical breakthroughs affecting one of these specific diseases (longevity jumps), or due to the potential of systematic underestimating of LEs for individuals suffering from them.

The idea of introducing the QxX family of indexes was perfectly logical; the availability of the objective indexes makes possible the creation of synthetic life settlement portfolios and the bringing in of new investors in the market who do not need to worry about the difficult-to-analyse risks that exist in the physical (as opposed to synthetic) life settlements market.

Some transactions using the index have been done. However, the timing

chosen for the index was most unfortunate, as shortly thereafter the financial crisis stopped the flow of capital to the life settlements market and very few transactions were being done even in the physical life settlements trading space. This also happened to be a time of great uncertainty for the life settlements investors, since the existence of systematic underestimating of LEs – underestimating of longevity risk – was becoming apparent with the announcements by some of the leading (in terms of market share but not always in term of expertise) LE providers that they were changing their methodologies to account for life expectancy greater than they previously anticipated. With the majority of investors affected by the uncertainty represented by this development, the market effectively halted: the number of new life settlements went down and the so-called tertiary trading (see Chapter 12) was still slow. At the end of 2009, Goldman chose to walk away from the index due to the low level of trading activity. While other, nonpublic and mostly ad hoc indexes have been designed and have resulted in actual transactions in the life settlement space, the decision by Goldman to no longer support QxX has made it harder to transfer the longevity risk of life settlements and to help the overall growth of that particular market.

Of course, beneficial as it is to have a tradable index, it is not a prerequisite to creating synthetic life settlements instruments (such as through swaps linked to performance of a specific large life settlement pool) and to creating effective hedging instruments. Still, the decision by Goldman to no longer support QxX is a big blow to the market.

Securitisation of life settlements

Another way to transfer the risk of longevity to a broad array of capital markets investors is through life settlements securitisation. In a portfolio of life settlements, the risk already resides in the capital markets. The transfer of the risk in the form of a securitisation opens the market to new investors and allows existing investors to free up their capital to buy more policies.

True public securitisation and true securitisation in general, with all the requirements such as that of true sale, is not the most likely path along which this market will develop, even if some such transactions are executed. On the other hand, the “weaker” form of securitisation, that of monetising the value of life settlement portfolios in private transactions, is likely to see some growth. Projections are very difficult given the unique and very strong challenges faced by the life settlement markets as a whole, on the one hand, and the very large potential market size on the other.

TRENDS AND EXPECTATIONS

Longevity risk has long been underestimated and neglected. The growing realisation of its magnitude and potential impact on DB pension plans, annuity providers, investors in life settlements and other holders of longevity risk are driving the search for the best solutions to transfer this risk to the capital markets.

Investors, for their part, also have interest in longevity risk being transferred to the capital markets, as it provides potential exposure to a weakly correlated risk factor and the resultant exotic beta. The market currently is rather inefficient in its early stages of development, presenting the more sophisticated investors with potential opportunities to generate greater return.

With the above general observations in mind, the following more specific trends and expectations apply when considering the issue of investment in longevity risk.

- ❑ We are witnessing the rapid transformation of the way longevity risk is thought of, along with growing (and justifiable) concern that continuing longevity improvements may lead to significant negative financial consequences for the holders of this risk.
- ❑ Longevity risk holders – the parties who are short longevity (or long longevity risk) – are becoming increasingly aware not only of the magnitude of the potential losses but also of the availability of the hedging tools for mitigating this risk, primarily through its transfer to the capital markets. A significant amount of innovation has taken place in the development of attractive new products for such risk transfer, with the focus on both providing high degree of hedge effectiveness and making the products attractive to investors.
- ❑ The process of identifying longevity risk is still in its early stages. Ways to properly quantify longevity risk are continuing to develop, helping defined benefit pension plans and other longevity risk holders to better understand and properly evaluate this risk.
- ❑ A necessary component of this process is the education of longevity risk holders, many of whom are unfamiliar with this type of risk, the way it might affect their assets and liabilities, and the magnitude of the risk they are already holding.
- ❑ While a number of instruments for the transfer of longevity risk to the capital markets have been developed, it is still unclear which of them will come to dominate the market and which will not be used. The outcome

could have a significant impact on the speed and direction of the market development. Greater standardisation will likely contribute to market growth, but only if the transferors of the risk are satisfied with the degree of hedge effectiveness achieved. It is possible that completely new structures, some of which have been discussed in general but never implemented, will emerge to supplement the existing array of longevity risk transfer tools.

- ❑ Over the past several years, significant advances have been made in modelling mortality and longevity. Stochastic mortality modelling is key to proper quantification of longevity risk; and, despite the progress already made, there is a need for better modelling tools to serve the interests of both longevity risk holders and potential investors in longevity insurance-linked securities.
- ❑ The development of new modelling tools requires better understanding of the drivers of longevity improvements and the factors important for longer-term stochastic projections of mortality and longevity. Continuing research is needed to improve the understanding of mortality dynamics, quantify the risk of longevity improvements more accurately and develop better pricing tools for longevity risk in its transfer to the capital markets.
- ❑ Further development of mortality and longevity indexes is required to provide better reference points for use in structuring longevity hedges and developing a market for longevity-linked securities. The indexes do not necessarily have to be used directly but can serve as building blocks for the constructing of longevity risk transfer instruments with high hedge effectiveness. Timely access to detailed and reliable population data is needed to construct better indexes.
- ❑ Similar to the need to educate longevity risk holders about this type of risk and the tools for its transfer to the capital markets, there is a need to educate investors and develop the expertise in the investor community required for proper analysis of this risk. This education process is a prerequisite to the development of longevity markets.
- ❑ Investor expertise should include the ability to manage longevity risk on a portfolio basis. While it is likely that investors almost always will be net long longevity risk (short longevity), portfolio management and optimisation tools can allow the risk to be reduced and greater risk-adjusted return to be generated.
- ❑ Longevity has always been the primary risk in life settlement investments. Gross underestimation of life expectancies has plagued the life settlements industry for years. In the analysis of these securities, the risk

is not so much in potential longevity improvements as in the basic misestimating of mortality rates applicable to the populations of insured individuals who have chosen to settle their policies. However, since the average life expectancy in life settlements is expected to increase, longevity improvements will play a greater role in pricing and actual investment performance of these investments.

- ❑ Portfolio hedging tools for life settlements are likely to be used more widely to manage the extension risk. Indexes used for this purpose are likely to be based on the longevity experience of life settlement pools as opposed to any proxy population.
- ❑ Reverse mortgages can also see the implementation of longevity hedges to protect lenders or providers of non-recourse loans from the risk of longevity being greater than anticipated.
- ❑ The types of investors that will most likely be willing to assume the risk of longevity – for proper compensation – include dedicated ILS funds and other hedge funds that do not have pension plans as their investors, life insurance companies, some family offices, and endowments. Other types of investors may become interested in this asset class as well.
- ❑ Central governments are holders of extremely large longevity risk. There have been calls on governments to issue longevity bonds or to pass their longevity risk to the capital markets in another way. But it is the size of the risk that makes it unlikely for any such measure to solve the problem. Capital markets are unlikely to be willing and able to assume the longevity risk that now resides with central governments in most countries. However, some governmental entities and local governments will likely make use of the capital markets solution to hedge at least some of their longevity risk.
- ❑ New regulations can become a catalyst for the rapid growth of the longevity markets. The current activity in the UK, albeit still limited, is based primarily on the regulatory actions that have forced pension plan sponsors and trustees to pay closer attention to such risks as longevity. Stricter regulations governing defined benefit pension plans can create an instant supply of longevity risk waiting to be transferred to longevity insurers or directly to the capital markets.
- ❑ The Life and Longevity Markets Association formed in the beginning of 2010 is focused on promoting standardisation in longevity risk transfer to the capital markets, with the goal of creating an active market in longevity trading.

After several years in which there were no actual transactions, the field of longevity risk transfer is now rapidly evolving. Most of the developments are happening in the UK, where the regulatory developments have provided the stimulus for quantifying longevity risk and transferring it to the capital markets. The longevity market appears to be ready to finally take off and to reach the critical mass needed to enable faster growth. Predictions of future growth of new markets are always dangerous: some types of insurance-linked securities saw initial growth but ended up disappearing almost without a trace. Longevity risk transfer to the capital markets – regardless of the form in which it is performed – may be an exception, however, as the market in all likelihood will continue its growth. The transfer of longevity risk to the capital markets addresses a need that will not be going away, and we are likely to see both continuing innovation and growth in market size over the next several years.

- 1 While in some countries all annuity products fall into this category, in others there are annuity types for which longevity does not present a risk.
- 2 A reverse mortgage is a loan made to a homeowner (typically a senior) against the equity in their home. The homeowner receives monthly or annual payments from the lender. The obligation to pay back the loan is deferred until the owner moves out or dies.
- 3 Based on Fitzpatrick (2009) and other presentations by Pension Corporation. Publications by the Pension Institute have included an estimate of 3% rather than 3.5%. Pension Corporation has also used the 3% figure in some of its presentations.
- 4 It appears unlikely that the longevity of the general population in the developed countries would decrease. Rather, actual longevity might end up being lower than that based on overly optimistic projections of continuing longevity improvements. In a small population such as that of pension plan members, longevity can also be lower than expected simply due to statistical fluctuations or misestimating of the mean.
- 5 While the idea of mortality forwards and similar instruments did not originate at JP Morgan, the LifeMetrics team deserves full credit for its development.
- 6 The currency swap in this specific transaction was supposed to be done to satisfy specific legal requirements to which the European Investment Bank is exposed.
- 7 In reality, there was a delay in the structuring of the bond, and, while the cohort of year 2003 was used, BNP Paribas attempted to place the bond only at the end of 2004.
- 8 The index was based on the actual longevity experience of the cohort of English and Welsh male population aged 65 years as published annually by the Office for National Statistics.
- 9 A more traditional solution for increasing capacity is for insurance and reinsurance companies to issue equity. In this case, while potentially a partial solution, it is probably not as cost-effective as using longevity bonds or longevity derivatives.
- 10 LifeMetrics tools and methodology have been developed by JP Morgan with the assistance of the Pension Institute. They are transparent and available on the JP Morgan website and at www.lifemetrics.com.
- 11 There is no standard terminology when it comes to being long (or short) longevity and longevity risk. Contradictory definitions are widely used. In this chapter, being long (that is, being the holder of) longevity means benefiting from greater-than-expected longevity; being short longevity means suffering negative financial consequences from greater longevity. Being long (that is, being the holder of) longevity risk has the opposite meaning to that of

being long longevity. It means being exposed to the negative effects of the risk of greater-than-expected longevity. In other words, being long longevity means being short longevity risk; being short longevity means being long longevity risk. Some analysts do not differentiate between the two usages; for them, being long longevity risk and being long longevity are synonymous. While both definitions have a certain internal logic, the use of these different terms interchangeably, imputing to them the same meaning, can introduce confusion.

- 12 JP Morgan has offered to transfer all rights to the LifeMetrics indexes to the LLMA.
- 13 Such a “contaminated source” might be a life settlements provider that has STOLI exposure (see Chapter 12) or a relationship with unscrupulous insurance agents that might guide policyholders on how to present their medical conditions so that their life expectancy is understated.

Part V

Managing Portfolios of Insurance Risk

Managing Portfolios of Catastrophe Risk

This chapter provides an overview of the topic of portfolio management of catastrophe risk, with a focus on catastrophe insurance-linked securities. Catastrophe reinsurance is also discussed, as well as some standard approaches to constructing and optimising a portfolio of catastrophe risk. The discussion is not limited to portfolios of property catastrophe risk but includes those that incorporate other catastrophe ILS, in particular securities that include the risk of catastrophic changes in mortality rates. While the technical details of the described approaches are outside the scope of this chapter, all of the main concepts are introduced.

Given that some practitioners have very strong insurance or reinsurance underwriting background but may lack the finance foundation, some of the basic finance concepts are introduced along the way.

PORTFOLIO CONSTRUCTION

Investments in risk, as we have defined them in Chapter 1, as well as the risks themselves, are always managed on a portfolio basis. While it is important to analyse each individual investment or risk carefully, ultimately it is the portfolio performance that matters. Individual components, important as they are, are relevant only in the context of their contribution to portfolio performance and overall risk.

The portfolio approach is equally important in investment and in insurance or reinsurance. Investors always want to optimise their portfolios, where optimisation is usually defined in terms of achieving the highest return for a chosen level of risk. Alternatively, portfolio optimisation could mean minimising the risk level for a given level of return. Optimisation is relevant to any investment portfolio, whether it contains stocks, bonds, commodity futures or catastrophe insurance-linked securities.

It is significant that measures of return and risk are not specified in the definition of the optimal or efficient portfolio. While the most commonly

used framework is that of mean-variance optimisation, return does not necessarily have to be defined as its expected value (mean), and risk does not have to be defined as the standard deviation of return. The mean-variance framework, while still widely used in investing, is slowly losing ground to the more sophisticated approaches. However, it does have some important advantages and provides useful insights, although even in a straightforward equity long-only strategy – if any strategy can ever be called straightforward – there are significant advantages to using other frameworks in addition to the mean-variance one. For many investment strategies and types of assets, measuring risk as simply the standard deviation of returns is not logical and could lead to investment losses.

Under modern portfolio theory (MPT), we wish to construct an investment portfolio that maximises reward and minimises risk by assuming that returns are represented by a normally distributed random variable, risk is measured by the standard deviation of returns, markets are efficient and investors behave in a rational manner. In its basic form, MPT assumes that the risk–reward preferences of an investor can be described by a quadratic utility function, and therefore only the mean and the variance of returns are important to the investor.

Panel 16.1, in very basic and inexact terms, describes the concept of the Markowitz-efficient frontier for those who come from (re)insurance background and need a reminder of these fundamental concepts. The efficient frontier in general, not only the Markowitz-efficient frontier under MPT, is convex,¹ which is a result of nonlinear changes in the risk–reward relationship corresponding to changes in the weights of individual components of a portfolio.

MPT provides a mathematical basis for diversification that can be obtained by assembling a portfolio of assets. Panel 16.2 shows, in very basic terms, how the risk measure used in MPT, the variance, can be reduced and is generally dependent on the degree of correlation among the assets in the portfolio. Following this simplified framework, we can see that there is a limit to diversification benefits, and that this limit is determined by the pairwise correlations among the components of the portfolio.


The limits on diversification are even clearer under the capital asset pricing model (CAPM). Arbitrage pricing theory, of which CAPM is a special case, leads to similar conclusions about diversification. Post-modern portfolio theory, a generalisation of MPT, leads to similar conclusions. These conclusions are all based on assuming, directly or implicitly, a degree of market efficiency and rational investor behaviour.

PANEL 16.1 SOME BASIC FINANCE CONCEPTS: THE MARKOWITZ EFFICIENT FRONTIER

Insurance and reinsurance practitioners sometimes have a very strong understanding of their field but do not have a clear picture of the fundamental finance concepts. For their benefit, some of these concepts are introduced here.

The mean–variance framework, where the measure of risk is the standard deviation (or variance) of returns and the measure of return is its expected value, results in a set of portfolios each of which has the highest level of expected return (reward) for a given standard deviation of returns (risk). Such a set, labelled the efficient frontier, is schematically shown in Figure 16.1. The return and its standard deviation (or variance) refer to the total portfolio, as opposed to an individual security it contains.

Figure 16.1 The Markowitz efficient frontier



The points on the efficient frontier represent the investment portfolios offering the highest reward for a given level of risk; the choice of a portfolio below the efficient frontier line would not be optimal. In this simplified framework, the primary goal of portfolio management becomes identifying and investing in the portfolio that (1) lies on the efficient frontier such as the one shown in Figure 16.1 and (2) carries no more risk than the investor is willing to assume. The second condition is investor-specific and depends on the investor's utility function.

PANEL 16.2 SOME BASIC FINANCE CONCEPTS: DIVERSIFICATION AND THE VARIANCE OF PORTFOLIO RETURN

The rate of return for a portfolio, R_p , can be defined as

$$R_p = \sum_{i=1}^N w_i R_i$$

where R_i is the return on asset i and w_i is the weight of this asset in the portfolio. The weights add up to 100%.

The variance of the returns for the portfolio can then be written as

$$\sigma_p^2 = \underbrace{\sum_{i=1}^N w_i^2 \sigma_i^2}_{N \text{ terms}} + \underbrace{\sum_{i=1}^N \sum_{\substack{j=1 \\ i \neq j}}^N w_i w_j \sigma_{ij}}_{N(N-1) \text{ terms}}$$

where σ_i is the variance of returns for asset i and σ_{ij} is the covariance between the returns on asset i and asset j . In a portfolio of many assets that are not too different from each other and whose weights are of the same order of magnitude, portfolio variance is approximately equal the average pair-wise covariance for the assets.

We can see that when there is a relatively large number of assets in the portfolio and none of them have considerably disproportionate weights, variance of the portfolio is typically dominated by the covariance term, while the variances of individual securities play only a minor role in their contribution to the portfolio variance.

In simple terms, this demonstrates reduction in volatility that results from diversification.

As a reminder to those (re)insurance practitioners who have not been exposed to finance since their college days, some of the basic concepts used in portfolio construction, optimisation, and performance measurement are described in simple terms in Panel 16.3. Beta combines the measures of correlation and volatility (standard deviation). It can be calculated, as in Panel 16.3, for an individual asset relative to a portfolio or to the whole market, or for a portfolio relative to the market. It can be used as a pricing tool under CAPM or as an important input in constructing the optimal portfolio.

Beta can be seen as a measure of systematic return. Alpha, the way it is defined in Panel 16.3, is the measure of idiosyncratic (unsystematic or specific) return, which is not dependent on the market movements. Alpha (again, only as it is defined in Panel 16.3), should have the expected value of zero based on CAPM. In practice, it is used as a measure of skill-based

investment return, which, on the expected basis, is positive for the “better” managers and negative for the “worse” ones.

Many asset managers, pointing out that the assumptions of market efficiency and the rationality of investor decisions in reality do not always hold, claim that they add value by generating positive alpha. Undeniably, skill makes a difference, especially in markets that are less efficient. Not all managers are created equal. We have to be careful, however, in attributing skill to managers who have shown positive alpha, even over a period of many years. It is important to understand the types of risk these managers are taking on (not limiting ourselves to measuring risk as the standard deviation of past returns) and the statistical flukes that sometimes lead to outperformance. There seems to be a clear reversion to the mean for most of the “better” managers as their track record grows longer.

PANEL 16.3 SOME BASIC FINANCE CONCEPTS: BETA, ALPHA AND THE SHARPE RATIO

As a reminder, the beta (β) of asset i in a portfolio is defined as

$$\beta_i = \frac{\sigma_{ip}}{\sigma_p^2}$$

where σ_{ip} is the covariance between the returns on asset i and on the portfolio. In the CAPM, the portfolio is the market portfolio. We can see beta as the result of linear regression analysis.

As another reminder, following CAPM, we can write for the market portfolio

$$E[R_i] - R_f = \beta_i (E[R_m] - R_f)$$

where R_m is the market return and R_f is the risk-free rate. Similarly, for a non-market portfolio p we can write

$$E[R_p] - R_f = \beta_p (E[R_m] - R_f)$$

Performing regression analysis over a certain time period t , we can write

$$R_{p,t} - R_f = \beta_{p,t} (R_{m,t} - R_f) + \alpha_{p,t}$$

In this formulation, beta is the slope of the regression line, while alpha shows how much better the portfolio performed relative to the expectation based on CAPM.

Sharpe ratio of portfolio p for the time period t , a related parameter, is defined as

$$\text{Sharpe ratio}_{p,t} = \frac{R_{p,t} - R_f}{\sigma_{p,t}}$$

EXOTIC BETA

Since insurance-linked securities as an asset class could be seen as a source of exotic beta for asset allocators, proper construction of a portfolio of ILS should not detract from this advantage. The “pure” insurance exposure is easiest to obtain in catastrophe ILS, since they are generally the ones with lowest correlation to the rest of the markets. And most of the ILS market is so inefficient that sometimes it may be relatively easy to generate abnormal returns on these investments.

A portfolio of catastrophe insurance risk is a source of exotic beta to investors, in the sense of providing return derived from exposure to an uncorrelated risk factor common to the asset class. Exotic beta, as the term is being used here, is different from the alternative beta defined as simply the beta from hedge fund exposure (hedge fund replication).

The ILS market is still evolving and is quite inefficient; so it represents a source of excess return to investors, offering exposure to a risk factor with return expectation above the “equilibrium” (efficient markets) level and low correlation with the global markets. In other words, this excess return results from the exotic beta qualities of the asset class in general. This asset class is particularly attractive due to its appeal as a source of exotic beta. Exotic beta, unlike traditional beta, is really nothing but alpha, as any positive excess return to a risk uncorrelated with the global market portfolio is alpha.² (For an investor having superior expertise in ILS, there is also the potential of generating additional, skill-based alpha, besides the exotic beta due to the exposure to this asset class.)

A portfolio of catastrophe insurance-linked securities can thus become a valuable alpha generator for investors. This advantage will not continue forever, though, since market inefficiencies always correct themselves; the exotic beta premium associated with catastrophe ILS and ILS in general will diminish and eventually disappear. Right now, however, that moment does not seem to be in the near future, since the current market inefficiency, to a significant degree, stems from insufficient investor education and a lack of expertise on the part of investors in the analysis of insurance risk and the management of ILS investment portfolios, and developing this expertise takes time. Other reasons for the inefficiencies vary by ILS market subsegment.

For an ILS portfolio manager, it is important to provide investors – at least those who are interested in this characteristic of the asset class – with the low correlation to the traditional asset classes and the alpha resulting from the inefficiencies of this market. To do so requires maximising the presence in

the portfolio of the “pure” insurance exposure and minimising, at least to some degree, the other, more traditional, risks and correlations.

HOW CATASTROPHE RISK IS DIFFERENT

Catastrophe insurance risk, by its very nature, is different from that found in most other investments. The traditional primary measures of risk, based on volatility in the form of relatively small price fluctuations, are of less significance in the analysis of many catastrophe insurance-linked securities or traditional catastrophe reinsurance. Instead, here the focus is on the risk of true catastrophes. This risk may be reflected only to a small degree in the historical returns. While the risk is present and fully reflected in the probabilistic forward-looking return distributions, even there it would rarely be properly measured by the traditional volatility measures such as standard deviation. Therefore, we need to focus on the measures of risk in the tail of the probability distribution. This is of critical importance in the analysis of an individual catastrophe insurance-linked security or a reinsurance contract; it is of less but still critical importance in the portfolio analysis of these securities.

Traditional measures of risk can work relatively well only in the case of normal or at least symmetrical probability distributions. Moderate deviations from these conditions can be addressed, at least to some degree, by using downside measures of risk, some of which are described later in the section on performance measurement. Still, these measures are rarely forward-looking, as they are typically used for performance measurement and for trying to use past prices of securities or option underlyings, with some adjustments, to make conclusions about future performance. This approach is inapplicable in catastrophe risk since there may have been no catastrophic events in the observation period. The measures of tail risk, often used for risk control more than for true risk measurement in most traditional asset classes, move to the forefront of the portfolio-management process in the case of catastrophe insurance risk.

While the statements above are relatively obvious and it may be easy to criticise the suitability of the mean-variance framework for the analysis of catastrophe risk, the fact remains that many investors judge asset manager performance based exactly on the parameters derived from that framework. For an asset allocator such as a fund of funds (FoF) or a pension fund, the Sharpe ratio and the return volatility (defined as its standard deviation) may be quite important in driving asset allocation decisions.

MEASURES OF RETURN AND RISK

The importance of correctly identifying relevant measures of reward and risk is difficult to overestimate. These are the key inputs into formulating an investment strategy and the necessary ingredients in constructing and optimising an investment portfolio.

In any optimisation framework, choosing appropriate measures of return is not independent from choosing the proper measures of risk. The two should correspond to each other. The minimalist approach of the mean-variance optimisation uses only two parameters of the probability distribution of possible outcomes (returns). In a perfect world, a quantitative asset manager would want to see the whole probability distribution, including its dependence on the many parameters affecting investment portfolio performance. Based on this dependence, he would then choose the portfolio that has the “best” distribution. A step in this direction would be to use several measures of reward and risk – even simply specifying several points on the distribution. The simplistic view that such an approach provides, however, is often far removed from practical reality.

Measures of return

Before discussing measures of return, we have to answer the question of how in general to define return. It can be relatively easy in the case of an instrument such as cat bond or a fully collateralised reinsurance contract. In cases when the collateralisation is absent or is only partial, the definition is more complicated. For example, when entering into a derivative transaction, such as buying or selling an exchange-traded cat derivative, the probabilistic cashflow models by themselves can only provide a distribution of the internal rate of return (IRR). To move from the IRR to the actual return, we need to know, for example, the cost associated with providing additional cash in the future, in case the total margin (the sum of maintenance and variation margins) increases. In the case of non-collateralised (or partially collateralised, or collateralised only from a certain time point in the contract term) catastrophe reinsurance underwritten by a reinsurer, there arises the same issue of the cost of having the ability to provide this additional capital in the future if required. This cost differs from one entity to another; it is also affected by its expected future actions.

Calculating the probability distribution for portfolio returns (as opposed to returns on an individual security) has similar problems that need to be addressed. We might in this case frame the problem in terms of the amount of cash needed to be held at any moment, or the level of liquidity of the port-

folio holdings that might need to be sold on short notice to satisfy margin requirements.

It is somewhat easier to identify the measures of return than the measures of risk. At the first level of approximation, the focus is almost always on the expected return, $E(R_p)$. The time horizon chosen might differ, but the measure of return as its expected (or actually realised) value is almost always the primary measure, and very often the only one.

Another example of a measure of return would be the probability of achieving a certain level of return, $P(R_p \geq R_{MAR})$, where R_{MAR} is the minimal accepted return. R_{MAR} can be set at the level of relevant benchmark. There can be more than one level of R_{MAR} , each with its own probability of being achieved. This could be seen as a goal or as a constraint. It could also be seen as a risk measure, since the complement of this probability is the risk of the return being below the specified level.

In the following discussion of risk measures, it becomes even more clear that return and risk cannot be considered in isolation, as the general goal is generating high risk-adjusted returns.

Measures of risk

There are two main types of risk and corresponding risk measures. One of them has to do with the volatility of returns. Standard deviation is one of the good measures of this risk, and that is the reason it is used in the mean-variance framework for portfolio construction and optimisation. Then, there are risks and corresponding risk measures dealing not with daily, monthly or quarterly volatility, but with catastrophic events; these are rare but could be devastating. Such events are likely absent from the historical period used for calculating the volatility risk measures. Disregarding their potential impact, however, is a risk not worth taking.

Below we discuss these two types of risk – the risk of relatively minor fluctuations and that of very large losses – and their corresponding risk measures, in a way that is somewhat simplistic since in reality there is more to risk than the two extremes. The whole spectrum of risks between these two extremes is important, and even risk measures seemingly belonging to one of the extremes are in some ways interconnected with those on the other end of the spectrum.

Volatility-related risk measures

Many of the risk measures dealing with volatility are so closely linked to returns that they might be more properly classified as measures of return, with the return being expressed on a risk-adjusted basis. Volatility in this

context can be defined in a number of ways beyond standard deviation or variance.

While this approach is not fully applicable to portfolios of catastrophe insurance-linked securities, the measures themselves are still important. Some are likely to be part of the reports to investors who are used to paying close attention to volatility and see it as an important, and often the most important, measure of risk. In fact, volatility-related measures play a role also in the catastrophe risk context, without regard to the need to demonstrate them to some investors. Especially as we move towards “less catastrophic” risks along the probability distribution, as opposed to those that hit a portfolio only once every few years, volatility becomes a more meaningful and significant risk measure.

Historical price volatility may not be as important a risk measure for catastrophe ILS as for other investments, but it is still a valuable determinant of risk. It can sometimes even measure true catastrophe risk, such as when it corresponds to a change in the market view towards tail risk that is not typically reflected in simple volatility measures. For example, recalibration of catastrophe models of the kind that followed the Katrina hurricane season in 2005 has the potential to affect market value of a portfolio. Alternatively, the environment might change, affecting the probabilities of catastrophic events tied to the components of the portfolio, changing their value and the value of the portfolio. An example can be developments that increase the chance of a pandemic leading to a catastrophic jump in mortality rates. The opposite example, that of the risk going down, is also valid and would lead to the portfolio value increasing. Actual losses in the portfolio, as long as they are isolated and not widespread, can lead to the same result.

Volatility can be a result of events in the rest of the financial markets, such as the effective dumping of catastrophe bonds by multi-strategy hedge funds in the second half of 2008 in order to generate cash to meet redemptions. As a result, the value of the bonds temporarily went down on the mark-to-market basis. The whole universe of the property cat bonds was affected by this phenomenon, with sudden correlation – both inter-portfolio and that with the rest of the financial markets – unexpectedly showing up. Some did not see this type of volatility as particularly relevant, since it did not change the probabilities of default, and had effect “only on the marks” and not the ultimate hold-to-maturity performance. But it did make a difference to any portfolio manager who reports results to investors (which is done on a mark-to-market basis) or who manages his portfolio on a more active basis than employing a simple buy (at issue) and hold strategy.

Some variation in the portfolio value is expected. Certain catastrophe risks exhibit seasonality affecting the market value of the portfolio. The risk of a North Atlantic hurricane making a landfall in the US is one example. Second-event catastrophe instruments have an even more pronounced seasonal behaviour.³ Second-event instruments that cover a period of more than one year are likely to increase in value if there has not been a qualified “first” event in the first year.

Measures of the time-bomb risk

Catastrophe risk, by its very definition, is the risk of very large losses. These losses are expected to happen only rarely. Any measure of this risk has to focus on the tail of the distribution of possible outcomes.

Value-at-risk (VaR) is defined as the maximum potential loss that can be incurred by the portfolio over a specified time period at a certain confidence level. It is the threshold value reached by portfolio loss over a specified time horizon at a given probability level.

If we define, following conventional notation, the cumulative distribution function (CDF) of portfolio returns X as $F_X(x) = P(X \leq x)$, then VaR at the confidence level of $1-\alpha$ is

$$VaR_{1-\alpha} = -\inf \left[x | F_X(x) \geq \alpha \right]$$

It is easy to see that the probable maximum loss (PML) measure commonly used in underwriting property catastrophe (re)insurance is a specific case of VaR. The VaR concept can also be used for non-tail (not catastrophic) events; for example, we might choose as a measure of risk and return the probability of portfolio returns being below (or exceeding) a certain level or benchmark.

Tail value-at-risk (TVaR) is the expected loss in the region of losses exceeding VaR. It allows us to see “beyond VaR” into the region of very large losses. VaR is the loss at the cut-off point beyond which the distribution of losses is not considered. TVaR, on the other hand, differentiates between two portfolios with the same VaR but different loss probabilities past that threshold. At the confidence level $1-\alpha$, TVaR is then

$$TVaR_{1-\alpha} = -E \left[X | X \leq -VaR_{1-\alpha} \right]$$

It can also be written in the following form

$$TVaR_{1-\alpha} = -\frac{1}{\alpha} \int_{-\infty}^{-VaR_{1-\alpha}} xf(x)dx$$

In other words, TVaR is the average of the worst $\alpha\%$ of possible outcomes. We can also write this in another way

$$TVaR_{1-\alpha} = \frac{1}{\alpha} \int_0^{\alpha} VaR_{1-\beta} d\beta$$

TVaR is often called conditional value-at-risk (CVaR) or tail conditional expectation (TCE), and the CVaR term is used as often as TVaR. (Depending on the definition used, TVaR and TCE can differ for distributions that are not continuous.⁴) We use the terms interchangeably, but define TVaR at the points of discontinuity as an average, in some cases weighted average, of the $TVaR^{high}$ and $TVaR^{low}$. The two are defined the following way: $TVaR_{1-\alpha}^{high} = -E[X | X < -VaR_{1-\alpha}]$ and $TVaR_{1-\alpha}^{low} = -E[X | X < -VaR_{1-\alpha}]$.

The terminology has not yet become standardised; so what some refer to as 99% VaR, others call 1% VaR (which is the difference between using α and $1-\alpha$); but the meaning is clear in either case and does not lead to confusion.

The random variable used in the definitions of VaR and TVaR is the portfolio profit, which is why the negative signs are used in the definitions to make sure VaR and TVaR are positive. There is no need for negative signs in the definitions if the value of losses is used instead, as is often done.

TVaR might be seen as more conservative than VaR, since it is always greater ($VaR_{1-\alpha} \leq TVaR_{1-\alpha}$). For some confidence levels α , TVaR may show a loss while VaR shows a gain on the portfolio ($VaR_{1-\alpha} \leq 0$ while $TVaR_{1-\alpha} \geq 0$). TVaR has some important advantages described below, but it is more difficult to interpret than VaR, making it a less useful measure in this regard.

Coherent risk measures

It has been advocated that a risk measure satisfy certain conditions that make it “coherent”. In part, the emphasis on coherent measures of risk stems from the criticism of VaR, which is widely used in risk management but has some undesirable properties. VaR is not a coherent risk measure, and the risk measures that are coherent overcome some of the problems with VaR. But this does not mean that they are always better and VaR should be abandoned.

The four properties of risk measure “coherency” – monotonicity, subadditivity, positive homogeneity and translation invariance – are detailed in Panel 16.4. They appear to be based on common sense, even though we can always find a potential problem even with such clearly defined and seemingly obvious properties. (For example, transaction costs and liquidity considerations might change when two portfolios are combined or when the size of an investment portfolio changes.)

VaR does not have the properties of a coherent risk measure. In particular, VaR is not subadditive. This by no means indicates that there is anything wrong with using this risk measure. It does mean that the results have to be carefully interpreted, and that the use of additional risk measures is likely to add value. It is possible to manipulate VaR by artificially reducing it at the expense of increasing the downside risk in the region beyond the VaR (often at the same time changing the expected return).

TVaR, on the other hand, is a coherent risk measure, as it satisfies the four conditions outlined in Panel 16.4. As such, it can be seen as a more logical risk measure to use. It certainly possesses the mathematical properties desired of a risk measure.

PANEL 16.4 COHERENT RISK MEASURES DEFINED

A coherent risk measure ρ is defined as one possessing the following properties:

1. Monotonicity: If $X_1 \leq X_2$ for all X_1 and X_2 , then $\rho(X_1) \leq \rho(X_2)$
2. Subadditivity: For all X_1 and X_2 , $\rho(X_1 + X_2) \leq \rho(X_1) + \rho(X_2)$
3. Positive homogeneity: For all X and all $\lambda \geq 0$, $\rho(\lambda X) = \lambda \rho(X)$
4. Translation invariance: For all X and α , $\rho(X + \alpha) = \rho(X) - \alpha$

The X above is the value of the portfolio.

A simplified explanation of the above four conditions is as follows. Monotonicity means that greater profits are associated with greater risks. Subadditivity means that a merger of two portfolios does not create extra risk, which is reflective of the concept of diversification. Positive homogeneity means that if the same portfolio is doubled or tripled in size, its risk will also double or triple. Translation invariance, in very basic terms, means that adding cash to a portfolio reduces its risk by the amount added.

Other comments on measures of risk and return

The critical difference between TVaR and VaR is that TVaR is subadditive, thus properly reflecting the concept of diversification. As discussed later, TVaR as a constraint makes the process of portfolio optimisation easier.

The time horizon chosen for calculating VaR and TVaR is of great significance. In traditional trading, the time horizon chosen in calculating these risk measures is typically very short, and daily VaRs are closely scrutinised. For portfolios of catastrophe risk, the proper time horizon is much longer, which somewhat changes the interpretation of these risk measures. This

does not mean that a trader should be at liberty to take unacceptable risks for a short period of time,⁵ such as massive use of highly leveraged derivative-type cat securities in the day before an expected hurricane landfall: it simply signifies that the portfolios of cat risk, with risk events expected to happen only rarely and the liquidity being limited, necessitates a longer view as opposed to the focus on very short time periods common in the trading environment. In fact, it is often beneficial to choose more than one time horizon in the calculation of these risk measures.

In the ideal world, all decisions would be made based on assessing the whole risk distribution as opposed to just one or two statistics derived from it. In practice, choosing several measures of risk and return is sufficient: it makes the decision-making process more transparent and intuitive, and it permits the use of portfolio optimisation techniques.

MANAGING A PORTFOLIO OF CAT RISK BY A (RE)INSURANCE COMPANY

The catastrophe risks – those of property insurance losses (from a hurricane, for instance) or life insurance losses (say, from a pandemic-related spike in mortality rates) – could be the same in an ILS investment portfolio and in an underwriting portfolio of an insurance or reinsurance company; but these portfolios are not managed the same way. A (re)insurance company never manages its portfolio of cat risks independently of the other facets of its operations. It faces both constraints and incentives that are different from those of an investor managing a portfolio of catastrophe risk – even if the investor's portfolio consists primarily of reinsurance-type instruments that can also be found in a cat risk portfolio of a reinsurance company.

For a (re)insurance company, there is always a trade-off between the cat risk it takes and the incremental return it generates on shareholder equity for the whole company. That return is a function of many variables, most of which are usually unrelated to the company's portfolio of cat risk. The cost of capital for a (re)insurance company is company-specific and plays a critical role in the decision of how much cat risk and at what price the company will take. The company has alternatives to assuming cat risk, including changing its underwriting volume distribution by line of insurance business, altering the risk profile of its portfolio by modifying its underwriting practices, or transferring (ceding) some of the cat risks in ways that are not available or cost-efficient for an investor managing a portfolio of cat risks. The investment portfolio of the company – its composition, investment returns, and relative riskiness – also has an effect on the way the company would want to construct its cat risk portfolio.

For most insurance and some reinsurance companies, the cat risk portfolio is not constructed in the traditional sense but rather is an outcome or by-product of the traditional, non-cat underwriting that results in the accumulation of cat risk in the general underwriting portfolio. For example, a large property insurance writer will likely end up holding significant catastrophe risk. This risk, comprising the cat risk “portfolio”, is seen as a necessary evil in writing traditional insurance. The company wants to minimise the risk and manage it very carefully. Often, some cat risk is passed on to a reinsurance company or to the capital markets if it is the most capital-efficient solution. For such companies, managing their cat risk portfolios is done indirectly through managing their overall, non-cat underwriting, and directly through ceding some of the unwanted risk to other parties. The pertinent decisions might be made only once a year; the rest of the time there is no cat portfolio management at all. In some cases, companies are more proactive and reassess their cat exposure more frequently, which could lead to buying additional reinsurance, changing their general underwriting, or entering into capital markets transactions such as an ILW or cat mortality swap.

In addition to the cost of capital mentioned above, an important and somewhat related consideration for (re)insurance companies is their financial strength ratings. In the US, the AM Best ratings are most important, while in the rest of the world S&P’s ratings are the ones most closely watched. This does not mean that ratings from Moody’s and Fitch do not carry weight. Each of these rating agencies has its own criteria and ways of handling cat risk in capital modelling and determining the minimum capital necessary to maintain a certain rating.⁶ Each insurance company has its own target rating; it is usually based on its marketing strategy and client base. In addition, there are certain rating thresholds below which a company can suffer negative effects such as losing policyholders or facing an automatic requirement to post collateral. These thresholds depend on the type of insurance sold by the companies, and on the jurisdictions involved. A company needs to maintain its ratings above these thresholds; it also has to be prepared for the consequences of an actual catastrophe. Such an event could lead not only to significant losses, but also to a downgrade below the threshold unless additional capital is quickly raised or the risk profile is altered. Keeping the probability of such a downgrade below a certain level is part of the process of managing the cat risk portfolio.

Risk-based capital (RBC) requirements imposed by regulators represent another important constraint. Not every jurisdiction has such requirements,

though. The US has an RBC framework, but the National Association of Insurance Commissioners (NAIC) RBC formula does not have an explicit charge related to the risks of natural catastrophes. This glaring absence of a critical risk will probably be corrected in the future, but this change can take time. The capital adequacy frameworks in Europe, both current and proposed, do generally take cat risk into account.

An issue related to the cost of capital is taxation. In some countries, insurance companies are allowed to establish a reserve to prefund loss payments for future catastrophes. In other countries (including the US) such reserves are not allowed to be recognised as a liability; this puts companies at a disadvantage in regard to taxation and makes it more expensive for them to assume catastrophe risk.

Professional cat underwriters

Some reinsurance companies do intentionally assume and manage catastrophe risk. There are some for whom underwriting catastrophe reinsurance is their main business. These are specialists who take extreme care in managing their cat portfolios to generate sufficient risk-adjusted return. Some of the most successful reinsurers focus primarily on underwriting cat risk. At the same time, some others have fallen into the category of the least successful precisely because of the cat risk they have underwritten.

These reinsurance companies are similar to investors in cat risk in terms of their thinking and the overall approach to managing cat risk, but the differences between them are still vast. Reinsurance companies are subject to all the constraints described above, such as the need to maintain certain ratings. Their cost-of-capital considerations are quite different from those faced by investors in insurance-linked securities.

Assuming reinsurance is essentially equivalent to the buy-and-hold strategy. You cannot get out of a reinsurance contract, and these contracts are not tradable instruments. This limitation leads to fewer active management options than are available to an investor in cat bonds. There is a great emphasis on properly constructing a cat risk portfolio, without making assumptions that there will be ways to get out of positions or make changes to the portfolio later on. This does not, however, mean that a reinsurer managing a portfolio of catastrophe risk does not have any hedging options. There is always retrocession, albeit often at a very high cost. There is an option of issuing a cat bond, but it can be expensive and time-consuming. There are options of entering into an ILW transaction or hedging the risk

using exchange-traded cat derivatives. Many of these options leave the reinsurer holding cat basis risk, but the high level of expertise often found at these specialist shops allows them to minimise this risk. Basis risk should never be neglected, and it is also important to remember that it can be heavily correlated with the model risk for the whole portfolio.

MANAGING A PORTFOLIO OF CATASTROPHE INSURANCE-LINKED SECURITIES

Managing a portfolio of catastrophe ILS is both similar to and different from managing a regular investment portfolio. While in some respects it is close to managing an underwriting portfolio of cat risk at a specialist reinsurance company, in other respects the differences are significant.

An investor does not have to deal with concerns such as RBC levels and company ratings that are important in the context of managing a portfolio of catastrophe risk by a (re)insurance company. The constraints that a reinsurance company has to deal with might be inconsistent with those resulting from internal economic capital modelling; but they still have to be considered. The task of managing an investment portfolio of insurance-linked securities appears to be “cleaner”, since many of these extraneous parameters do not need to be considered. That does not make it easier, though, only different.

The investor is exposed to the volatility of mark-to-market valuation, which in the absence of a catastrophe is not a concern for insurance and reinsurance companies that have a smooth pattern of earning premiums (recognising revenues).

Since an investor such as a dedicated ILS fund or a multi-strategy hedge fund might be hit by redemptions and the need to liquidate some of the investments, a certain level of liquidity should be maintained, and cash and liquidity management policy established and followed. While an insurance or reinsurance company can have liquidity concerns as well, they rarely trickle down and affect the way an underwriting portfolio of cat risk is managed.

The main steps in the ILS portfolio-management process are common to traditional asset management. It is a systematic process that continuously goes through predefined decision loops. The main steps are the following.

- Formulation of investment policy, often resulting in a formalised investment policy statement, is the first step in portfolio management. An investment policy includes, but is not limited to, investment goals and

general ways to achieve them as well as any legal or other requirements appropriate for an investor in insurance-linked securities. It likely includes the types of insurance-linked securities and other instruments that can be used to achieve the investment goals and may also contain some restrictions such as limits on the use of leverage. Concurrent with formulating an investment policy, and to a significant degree included in it, is the determination of return and risk objectives of the overall investment strategy. In the case of ILS, the choice and definition of the return and risk measures are of particular importance, as is the investment time horizon. Analysis of the market conditions and opportunities consistent with the risk and return objectives and the constraints then becomes the key input in formulating specific investment choices. The three main markets in the analysis are the broadly defined insurance-linked securities market, the global financial market, and the insurance/reinsurance market.

- ❑ Constructing the optimal portfolio for the chosen overall investment strategy is the next step in the process. It can involve strategic asset allocation by ILS type, followed by the security analysis combined with portfolio optimisation, which results in the security selection. The step of strategic asset allocation can be bypassed and the optimisation performed using the whole universe of available insurance-linked securities. The portfolio is constructed in a way that incorporates expectations related to the markets and individual securities, which are all taken into account in the optimisation. Strategic asset allocation, when performed, can also be a product of optimisation and often reflects the skill set of the portfolio managers and their ability to properly analyse different types of insurance-linked securities. The final step is that of execution: taking long or short positions in the selected securities, implementing additional strategies and maintaining liquidity and other rules appropriate for the portfolio.
- ❑ Managing the already constructed ILS portfolio is a dynamic process aimed at reoptimising the portfolio based on changes in the market environment and the portfolio itself, investor feedback and the results of modelling of individual securities. New opportunities, such as those presented in the secondary market, necessitate the analysis of whether portfolio changes are necessary to best meet the investment objectives. The process must be truly dynamic and proactive in order to take advantage of tactical opportunities such as those created by “live cat” trading. At the same time, the process involves closely monitoring potential risks

to the portfolio, whether they are related to individual holdings, changes in the general market environment, developments in the ILS space, occurrence of natural catastrophes or the revision of probabilities of cat events such as pandemic-flu-related mortality spikes. Any of the above might require portfolio rebalancing or other changes. The availability of new hedging tools, or price changes in existing tools, might expand the available options and lead to opportunities to achieve greater risk-adjusted return. The process of managing an ILS portfolio should also incorporate risk-management rules and procedures to minimise all types of risks to the portfolio, including operational risk. Portfolio optimisation tools are also tools of risk management, and should be used as such. As the environment changes and the portfolio changes with it, it might be necessary to develop new stress tests for the portfolio to reflect these changes. It is not enough to have a portfolio optimised based on its calculated return distribution: there should be specific policies and procedures in place to deal with situations when things go wrong – due to either the occurrence of a covered catastrophic event, or an external factor such as large-scale redemption requests. This too is part of the risk-management process, which in turn is part of the overall portfolio management process.

Risk management is of critical importance in ILS portfolio management. Rather than simply being a risk-control tool, it serves as an essential input in portfolio optimisation and thus affects portfolio composition. This process may reveal that an investment portfolio does not have a sufficient risk level, and additional risk needs to be taken to bring the portfolio closer to being optimal. It is essential that the risk manager be more than a risk cop and instead become part of the decision-making process.

INSTRUMENTS TYPES

The universe of investment instruments available to an ILS portfolio manager is an important determinant of the overall strategy and the optimisation techniques that can be used. In a general case of portfolio construction and optimisation, we would want to have as many options as possible, including access to the greatest number of security types and individual securities of each type.

The main types of catastrophe insurance-linked securities are:

- ☐ cat bonds;
- ☐ industry loss warranties (ILWs);

- ☐ catastrophe collateralised reinsurance;
- ☐ catastrophe derivatives (exchange-traded or over-the-counter);
- ☐ some sidecars;
- ☐ extreme mortality bonds;
- ☐ extreme mortality derivatives; and
- ☐ contingent capital notes.

Even though some types of cat ILS are more important than others in portfolio construction and management, being able to access as many of these types as possible, either as investments or as hedging tools, gives a portfolio manager flexibility and creates new options in portfolio optimisation. Right now, extreme mortality securities are usually not found in an ILS portfolio. The same can be said for contingent capital notes, especially since they are very uncommon nowadays. Not using extreme mortality securities is not a significant limitation for a manager of cat ILS. Not using the whole spectrum of property cat ILS is. Those managers who do not limit themselves to investing in cat bonds have added flexibility, since, as markets conditions change, they can redeploy capital to take advantage of the most promising opportunities. Dedicated ILS funds that are legally restricted to investing in cat bonds can find themselves at the mercy of the markets if the supply of new cat bond issues decreases, or if the pricing levels make cat bonds temporarily unattractive. Funds able to write collateralised reinsurance or at least ILWs may under some circumstances have an immediate advantage over those that are limited to cat bonds.

In reality, the situation is more complicated. While it is true that having access to as broad a universe of catastrophe ILS as possible creates new options and portfolio optimisation opportunities, this advantage comes with a cost of having to develop additional infrastructure and, more importantly, expertise in these other types of catastrophe insurance-linked securities. The question of having the expertise is critical, as there are many pitfalls for an unwary investor looking into a new type of ILS. This caveat applies particularly to collateralised reinsurance, which requires reinsurance expertise.

Those portfolio managers who do not run dedicated ILS funds with a singular focus on catastrophe ILS also have the option of investing in other, non-catastrophe types of insurance-linked securities; or deploying their capital in asset classes unrelated to insurance, and investing in cat or other ILS only when the pricing levels are at their highest.

PORTFOLIO CONSTRAINTS

Constraints used in the management of a portfolio of cat insurance-linked securities come from two primary sources. One has to do with the general mandate given to the portfolio manager: that is, the constraints specified in the fund documents of a dedicated ILS fund; or, when the ILS portfolio is really a sub-portfolio of a bigger asset portfolio such as that of a pension fund or a multi-strategy hedge fund, the constraints imposed by the manager of this bigger portfolio. Constraints driven by legal or tax considerations, such as those that have to do with the potential limit on insurance activities in some fund structures, are also in this category. The second type of constraint is imposed largely internally, in order to avoid excessive risk and to maintain a fund risk profile consistent with the chosen management style, risk appetite and implicit promises made to investors.

Constraints are important in risk management. They are also key elements of a portfolio-optimisation framework. VaR and TVaR are the constraints that can be used in portfolio management in general and are important in the management of catastrophe risk in particular.

While TVaR may have won over VaR based on theoretical considerations (in particular because it is a coherent risk measure), the same outcome has not yet occurred in the practical use of these risk measures. The VaR has already been embraced by many risk managers, especially in the banking industry, as well as some regulators. TVaR appeared later, when the VaR culture was already established. TVaR has been slowly but steadily gaining ground since then. Probably even more importantly, TVaR is more difficult to interpret and is seen by some as a less relevant measure than VaR. Asset managers are concerned with their losses not exceeding a certain threshold, of which VaR is a good measure, so it can be used as a constraint in the portfolio management process. Some managers have less concern about what happens beyond this point – “it is all lost anyway” – and thus see the TVaR measure as less relevant. They believe that losses should not reach that level, so that is the main and possibly only constraint related to tail risk. From that point of view, tail risk management simply means making sure that losses do not exceed this level. These managers see VaR not so much as a constraint but as a parameter that needs to be minimised (at least within a certain range of values). Some do not see TVaR as a relevant constraint for somewhat personal rather than purely business reasons, in terms of career risk: they believe that if losses reach the VaR level, investors in their funds will withdraw their money anyway; or they would probably lose their jobs if their portfolios were to suffer massive losses.

TVaR, despite its attractive mathematical properties, is a difficult constraint to specify for any manager. Is losing 55% of the portfolio averaged over the 1% of worst outcomes over a one-year time horizon the right constraint in constructing and optimising a portfolio? Or should it be 45% or 75% instead of 55%? Is $\alpha = 1\%$ (once every 100 years) the right choice, or should 0.4% (once every 250 years) be chosen in specifying the constraint? There are no certain answers to these questions. Choosing constraints involving VaR is easier and more intuitive, though even here there are no easy answers.⁷

In some cases, it makes sense to choose several tail-risk constraints, for example in the form of VaR at more than one confidence interval and/or more than one TVaR. (Later we examine how these choices affect the portfolio optimisation process).

Probably the most intuitive type of constraint is the requirement that the probability of returns being negative or below some benchmark not exceed a certain level. While such constraints do not usually deal with tail risk, it is easy to see that they are equivalent to using VaR as a constraint in the optimisation process.

An example of portfolio constraints is the following constraint set for an initial portfolio of US\$400 million:

$$\begin{cases} 95\% \text{ VaR} \leq \text{US\$50M} \\ 99\% \text{ VaR} \leq \text{US\$160M} \\ 99.9\% \text{ VaR} \leq \text{US\$350M} \\ 99\% \text{ TVaR} \leq \text{US\$250M} \\ P(R_p \leq 0) \leq 15\% \\ P(R_p \geq 9\%) \geq 70\% \end{cases}$$

where all risk-and-return measures are defined using a one-year time horizon. This is likely only a subset of the broader set of portfolio constraints.

In some optimisation frameworks, there are the two types of constraint:

- hard constraints, which have to be satisfied in the optimal decision; and
- soft constraints, which indicate preference of some solutions over others, expressed in the form of additional (local) objective functions.

Constraints used in ILS portfolio optimisation are not limited to those related to risk but also include those having to do with the availability of certain portfolio options (for example, limits on how much of a specific security can be bought, either in general or at a specific price), and general

commonsense constraints. The latter could include "safeguard" constraints such as the minimum number of securities and the maximum position size.

In a broader formal framework, we can speak about risk constraints $\rho_j(P) \leq u_j$, where $\rho_j(P)$ is the risk measure used in the j th constraint. (See also constraints in the linear programming problem in Panel 16.5.) If we consider all portfolio constraints, we can write M constraints $I_j \leq C_j \leq u_j$, where j takes values from 1 to M .

Most of these constraints are not unique to managing a portfolio of catastrophe ILS but are commonly used in the investment management of other types of assets. Constraints specific to insurance-linked securities are those that have to do with the ILS market inefficiency, limits on diversification within an ILS portfolio due not only to the small size of the market but also the small number of "risk buckets", and the unique properties of cat insurance risk. The general framework, however, is common to most asset classes.

STANDARD TOOLS AND THE MODELLING OF INDIVIDUAL SECURITIES

The standard tools used in cat risk portfolio modelling (as opposed to optimisation tools described later) are the catastrophe models provided by the three firms, AIR Worldwide, EQECAT and Risk Management Solutions (RMS). The basic structure of these models has been described in Chapter 3 and Chapter 4. They are used for modelling the risk of natural catastrophes and their impact on insured losses.

How such a model is used for the analysis of individual securities is described in the chapters that provide an overview of the models and catastrophe modelling process in general. The three modelling firms provide both the general insurance natural-catastrophe models and the models specifically designed for ILS investors (which are based on the general-catastrophe models). In addition to the three firms mentioned, there are a couple of competitors that have created catastrophe models for specific territories and perils. They are generally not used in the ILS analysis and used only rarely in the analysis of catastrophe insurance risk in general; the field is completely dominated by AIR Worldwide, EQECAT and RMS.

The models created for investors – CATRADER by AIR, eCAT by EQECAT and Miu by RMS – allow the analysis of catastrophe bonds, industry loss warranties and exchange-traded catastrophe derivatives in the portfolio context as well as individually. One of them can also be used to manage portfolios of catastrophe reinsurance, in addition to pure insurance-

linked securities. (The other two also provide this capability, at least to some extent, but only in an indirect way requiring the use of consulting services of the model provider. This immediately reduces their usefulness for those investors who want to include reinsurance contracts in their portfolios.)

Individual security analysis: cat bond

The analysis of a security such as a cat bond starts not with the use of a model but with careful examination of the investor documentation and the risk analysis included in the offering circular (OC). Figure 16.2 illustrates some of the considerations important in the analysis of a cat bond as well as the primary output of the analysis. The analytical framework outlined there has as its main focus the analysis of a cat bond on a standalone basis, before the security is included or considered for inclusion in an investment portfolio. The primary output of the analysis focuses on the loss distribution and its sensitivity; pricing considerations are part of the next stage of the analytical process.

The example in Figure 16.2 shows how the initial analysis concentrates on both quantitative (contained in the risk analysis section of the OC) and qualitative factors, with qualitative factors often playing as great a role as the quantitative ones. The goal is to convert as many of the qualitative factors as possible into quantitative inputs or adjustments, ranging from quantifiable information on the composition of the underwriting portfolio or the quality of the underwriting of the sponsor for indemnity transactions, to stress scenarios and types of sensitivity testing deemed reasonable based on qualitative considerations. An important consideration is whether the investor believes it has superior information to improve on and expand the analysis already presented in the OC. The investor might have particular insight into the degree of accuracy or biases of the model used for some of the peril/territory combinations; knowledge of the underwriting practices of the sponsor; or superior ability to identify and understand the non-modelled risks. It is such knowledge, and the ability to use it in the analysis, that results in competitive advantage.

Similar analysis can be performed for ILWs or exchange-traded derivatives such as IFEX contracts. Here, too, having an informational advantage and the ability to use it in the analytical framework can be of great importance. Of even greater importance is the ability to optimise a portfolio, which makes use of the same modelling tools.

In addition to the property catastrophe risk models discussed above, RMS has developed a pandemic model useful in analysing the risk of catastrophe

Figure 16.2 Considerations in the analysis of a cat bond outside of portfolio context

Starting point	Quantitative risk analysis in the offering circular (OC)	Quantitative information in the offering circular (OC)
Primary considerations in the analysis	<ul style="list-style-type: none"> – Type of default trigger – Peril(s) and territories covered – Modelling firm providing the analysis in the OC – Confidence in the model for this type of peril and territory – Types of modelling and sensitivity testing performed – Level of data detail provided – Short- and long-term hurricane forecasts – Securitisation structure – Bond tenor – Credit rating – Quality of collateral and swap counterparty rating – Legal, regulatory and tax considerations – Subordination level – Issuance amount – Identity of the sponsor / issuer 	
Additional considerations	<ul style="list-style-type: none"> – Secondary pricing levels – Reinsurance pricing levels for the risks embedded in the security – Input from rating agency analysts, modelling firms, strategic partners and other investors – Market intelligence from a variety of sources – Latest scientific research relevant to the risk involved – Whether the investor has superior knowledge of the specific risks embedded in the security – Possibility of model arbitrage – Foreign exchange risk – Degree of ease or difficulty of replicating the security using other ILS and price levels for such replication – Hedging instruments available for all or part of the risk embedded in the security; Cost of hedging; Potential basis risk 	

Primary output of the analysis

- Loss probability distribution (based on the investor analysis and not necessarily directly corresponding to the loss exceedance curve provided in the offering circular)
- Qualitatively defined degree of uncertainty and model risk underlying the loss distribution
- Main contributors to modelling risk
- Stress scenarios appropriate for the security given the modelling uncertainty
- Risk distribution by season and over the term of risk exposure
- View of the future liquidity of the security
- “Cleanness” of exit
- Map of risk exposure that could also be used to assess correlation with other securities in the portfolio and the ILS universe

Note: Remodelling data is not included in this initial analysis; it is described later in the context of portfolio modelling.

mortality (as discussed in Chapter 11). The output of this model is used in order to incorporate in Miu all outstanding extreme mortality bonds. Miu does not, however, have the ability to model extreme mortality risks other

than these bonds. (Tools provided by all three cat modelling firms constantly evolve; new features are added all the time.)

Risks not modelled using standard tools

Some catastrophe risks cannot be modelled using the standard tools discussed above. For example, a collateralised reinsurance contract covering catastrophic risks of aviation or satellite insurance does not lend itself to being modelled using these modelling tools. Investors having the expertise to analyse the risk would create their own models. The resulting loss distribution can then be used in the general portfolio optimisation process. This could be done in a relatively simple way, with the exception of the risks correlated with other securities such as cat bonds. This correlation cannot be easily incorporated in the modelling process used for portfolio optimisation.

PORTFOLIO OPTIMISATION

As discussed above, a portfolio manager should continually work on monitoring the portfolio and the market environment to take advantage of the new opportunities to accomplish the investment goals and to make sure that the risk levels remain consistent with the specified risk constraints. Managing a portfolio of catastrophe ILS and catastrophe insurance risks in general has the same goals; the difference comes only in the way they are accomplished.

The portfolio construction and optimisation process consists of four elements.

1. Identification of the universe of available instruments – types of catastrophe insurance-linked securities or reinsurance as well as the specific securities available – with the understanding that the latter can change constantly due to changes in secondary markets availability and pricing, as well as the specific execution or trading options available to the portfolio manager. This step has as its natural outcome the set of decision variables that can be used in the optimisation process. It also results in identification of some constraints that are used in portfolio optimisation.
2. Based on the analysis of risk and return preferences and goals, formulation of the objective function that will be maximised or minimised in the optimisation process. The most common objective function is the expected return on the portfolio over a certain period of time, which in most cases can be written as the linear combination of the returns on the portfolio components, or $R_p = \sum w_i R_i$, where w_i 's are security

weights in the portfolio. The process of optimisation may have more than one objective function corresponding to possible multiple objectives in the management of an investment portfolio. In some cases, multiple objectives can be reflected in one objective function, to a degree reflecting some trade-off between risk and return.

3. Formulation and formalisation of constraints, which can include constraints on tail risk in the form of VaR, TVaR or some other measures for a specific time horizon or set of time horizons; constraints having to do with the expected frequency of achieving (or falling below) certain return levels; constraints on the expected volatility or downside risk measures not directly involving tail risk; constraints imposed by practical limitations such as the level of available funds (and possible levels of borrowing, if any) or externally imposed liquidity requirements; constraints on the amounts of individual securities that can be bought or sold at specific prices; constraints on the ability to change several positions or implement several investment decisions simultaneously;⁸ and many others. The list of constraints can be very long, but the length of the list, while obviously having an effect on the time needed to run an optimisation programme, is less important for the process of optimisation than the types of constraints being used.
4. The last element of the process is running an optimiser – software that is based on optimisation algorithms – to identify the “optimal” portfolio for the given set of constraints and objective function(s), and to execute the trades or perform other actions to move from the current to the “optimal” portfolio. Then the process repeats itself, even though the ILS markets are very slow by today’s standards, and actual changes to the portfolio are infrequent for the vast majority of ILS investors.

The framework described above appears conceptually simple. Direct practical implementation, however, is impossible, even if we had an optimiser to handle the problem as formulated. Numerous qualitative judgement calls and decisions have to be made at every step of the process, and simplifying assumptions are unavoidable. In truth there is never a magic software model that takes in some inputs and spits out the correct actions for the portfolio manager to take. While the same statement can be made about managing a portfolio of any asset class, it is especially true in managing a portfolio of catastrophe insurance risk, where the risks and returns are far from obvious and, despite the seeming abundance of quantitative information, qualitative factors are of paramount importance.

Use of standard tools in optimisation

The standard cat modelling tools described above – CATRADER from AIR Worldwide, eCAT from EQECAT and Miu from Risk Management Solutions – allow an investor to map most cat ILS as a portfolio. All three come with libraries of modelled outstanding cat bonds that are provided to qualified investors.⁹ (See more on the cat bond analyses libraries in “Remodelling and portfolio optimisation”.)

We can see the distribution of outcomes for individual scenarios, where a scenario can be a catastrophic event or a simulated year. Such a model can simulate the results over a period of 100,000 years and much longer, resulting in a probability distribution of losses for a specific portfolio without the need to make assumptions about correlation among the securities. This data can be the basis for portfolio optimisation.

Several points have to be taken into account, though. An investor having a certain view on an individual insurance-linked security – such as believing that for a particular cat bond the loss distribution differs from that in the offering circular – has to make these adjustments before such a security is considered to be part of the ILS universe used in optimisation. This task is not always easy. Not all of the three software packages provide the ability to effectively model collateralised reinsurance (with the exception of industry loss warranties). Only one of the three – Miu from RMS – provides information on extreme-mortality bonds and allows the user to incorporate them in hypothetical portfolios. This limitation of the other two models can be overcome by analysing extreme-mortality securities outside of the model and then assuming zero correlation (or making another assumption) between extreme-mortality securities and those linked to the risk of natural catastrophes.

The data produced by the models and serving, with adjustments, as the basis for optimisation includes only information on loss distributions. Pricing information is not provided and is incorporated in the process later as another set of inputs for the portfolio optimiser.

Linear programming

A number of optimisation algorithms can be used in portfolio construction. The best known is linear programming, which can be used for optimisation problems such as the one outlined in Panel 16.5.

Linear programming methods are very convenient, in part due to their simplicity and relatively high computational speed. It is important to note that the TVaR constraint can be used in linear programming algorithms

PANEL 16.5 LINEAR PROGRAMMING

The standard formulation of the linear programming problem is maximisation of a variable F , which is a linear function of N non-negative decision variables x_i

$$F = \sum_{i=1}^N w_i x_i$$

subject to M constraints on linear combinations of decision variables

$$\sum_{i=1}^N g_{ji} x_i \leq u_j,$$

where j takes values from 1 to M .

In the matrix form, this can be written as maximisation of $\mathbf{W}^T \mathbf{X}$ subject to $\mathbf{GX} \leq \mathbf{U}$.

The same problem can be rewritten in terms of minimising the F variable and the constraints can also establish floors instead of ceilings on the linear combinations of decision variables. A constraint can also be an equality, where a linear combination of decision variables is constrained to equal a certain value. The linearity of both the variable being maximised (return) or minimised (risk) and the constraint variables allows the use of simple and usually efficient optimisation techniques of linear programming. An observation can be made that a maximum (or minimum) of the objective function exists in the corner of the constraint set since the function is linear. In some cases, it can also exist along the entire surface of the constraint set.

This standard formulation can be expanded in a number of ways to still be able to use the linear optimisation algorithms.

while a constraint based on regular VaR cannot. This is an important advantage of TVaR over VaR constraints in portfolio optimisation.

There are numerous optimisation software packages, many of which can be used for portfolio optimisation. Most are based on one of the linear programming algorithms. However, the number of approaches to optimisation and various algorithms is vast and certainly not limited to linear programming.

Formulating the optimisation problem

Formulating the catastrophe ILS portfolio optimisation problem is done the same way as for any other asset portfolio. All the key elements have been described above, including identification of the variables used in optimisa-

tion, formulation of the constraints and determination of the objective function. The specific constraints and, in some cases, the objective functions may be unique to ILS, but the framework is generic. To summarise, the problem is that of maximising or minimising objective function(s) $F(F_k)$ subject to the set of constraints $l_j \leq C_j \leq u_j$ (some of which are risk constraints $\rho_j(P) \leq u_j$), which is the standard constrained optimisation problem.

The objective function is often the expected portfolio profit or portfolio value at the end of the one-year time period; in other words, we want to maximise $E(R_p)$ or $E(P)$ in one year. TVaR at one or more levels of significance is an example of possible constraints. There are also constraints having to do with position sizes ($w_i^{low} \leq w_i \leq w_i^{high}$), determined in part by the security availability and in part by risk management policies. Defined this way, the optimisation problem can usually be converted to a form that allows the use of linear programming algorithms. If constraints such as those involving VaR or specifying certain thresholds on the return probability distribution are introduced, the standard linear programming approach is no longer applicable and other portfolio optimisation methods should be utilised.

Example of optimised portfolio statistics

The result of the optimisation process might be a portfolio with parameters such as those shown in Figure 16.3. For a hypothetical portfolio, it shows the expected annual return of 16% (before fees) with the standard deviation of returns of 14%. For a dedicated ILS fund with a 2-and-20 incentive compensation structure, this translates into the return of 11% to investors.¹⁰ A look at the several points taken from the probability distribution of returns suggests that the components of the portfolio are likely not simply cat bonds; rather, it is likely that the portfolio includes such instruments as collateralised reinsurance and ILWs.¹¹

The values of VaR and TVaR are instructive: annual portfolio losses are expected to be worse than 10% of the beginning portfolio value only once every 20 years, worse than 36% of the portfolio once every 100 years, and worse than 51% once every 250 years. Leaving aside the question of whether returns are adequate, which is dependent on the perspective of an individual investor, we can observe that the risk – as quantified by the three measures above – is not significant for a portfolio of catastrophe risk. We should not, however, limit the risk view to observing only these measures. A more careful examination of the whole return distribution and other risk measures might lead some investors to another conclusion, especially when the expected return is considered.

As expected, we can see that TVaR is greater than VaR for the same confidence level. The difference is smaller very far in the tail of the distribution but grows quickly as the more frequent events are considered.

The traditional way of presenting VaR measures is not as a percentage loss (negative return), as shown in Figure 16.3, but rather in absolute values such as the dollar amount of loss. If we assume that the hypothetical portfolio had US\$250 million at the beginning, then the Var@99\% of 36% translates into US\$90 million as shown in Figure 16.4 overleaf. As above, the risk measures in this hypothetical portfolio are calculated using the one-year time horizon.

The values of the expected return, standard deviation, VaR and TVaR at several levels of confidence, as well as probabilities of the return being non-negative and exceeding several specified levels, are calculated based on the probability distribution of portfolio returns. The return cumulative probability distribution that was used for this illustration is shown in Figure 16.5. We can again notice that the portfolio's risk-return profile appears to be rather attractive, and that it likely includes a significant amount of catastrophe reinsurance or similar instruments, as opposed to being purely a cat bond portfolio. The probability distribution for a cat bond portfolio is rarely

Figure 16.3 Some of the risk and return parameters of a hypothetical portfolio of insurance-linked securities

Hypothetical ILS portfolio	Expected return (<i>before fees</i>)							16%
	Standard deviation of return							14%

Probability of annual return being LESS than								
-40%	-30%	-20%	-10%	0%	10%	20%	30%	40%
0.8%	1.3%	2.4%	5.0%	9.6%	24.9%	58.3%	91.0%	99.9%

Value-at-Risk		Tail Value-at-Risk	
VaR@99.9%	61%	TVaR@99.9%	63%
VaR@99.6%	51%	TVaR@99.6%	58%
VaR@99%	36%	TVaR@99%	50%
VaR@98%	23%	TVaR@98%	44%
VaR@97%	17%	TVaR@97%	41%
VaR@96%	13%	TVaR@96%	39%
VaR@95%	10%	TVaR@95%	37%

Figure 16.4 VaR and TVaR for the hypothetical portfolio of insurance-linked securities expressed as dollar amounts following the standard convention

Hypothetical ILS portfolio		Initial portfolio value (in millions of US\$)	250
Value-at-Risk		Tail Value-at-Risk	
VaR@99.9%	153	TVaR@99.9%	157
VaR@99.6%	128	TVaR@99.6%	144
VaR@99%	90	TVaR@99%	125
VaR@98%	57	TVaR@98%	109
VaR@97%	42	TVaR@97%	101
VaR@96%	32	TVaR@96%	96
VaR@95%	25	TVaR@95%	92

that smooth and not always continuous, given that the number of cat bonds in it cannot be very large, the risks are usually concentrated in only a few risk buckets, and many of the cat bond outcomes are almost binary (since, if a cat bond defaults due to a cat event, there is a high probability of its being a full loss).


The probability distribution of returns for the whole portfolio is the output of the optimisation process that identifies the optimal portfolio. The sensitivity of this distribution to the input parameters has to be tested. Useful information can also be obtained in those cases where confidence intervals for parameters such as VaR can be estimated.

Marginal impact of investment options

The optimisation process can be seen as the consideration of all possible scenarios – that is, the various possible investment decisions and the resulting portfolios – and then choosing the scenario that results in the “best” portfolio that maximises or minimises the objective function(s) given the imposed constraints. (In reality, the optimisation algorithms usually allow us to avoid looking at all possible scenarios and instead leads to quicker convergence to the maximum or minimum.)

Figure 16.6 illustrates the basic schematics of the decision-making process when several options are presented. In this concept illustration, the focus is on the marginal impact of each option on the portfolio risk–return profile.¹² The framework, even in this most simplistic form chosen for illustrative

Figure 16.5 Cumulative probability distribution of annual returns for the hypothetical portfolio

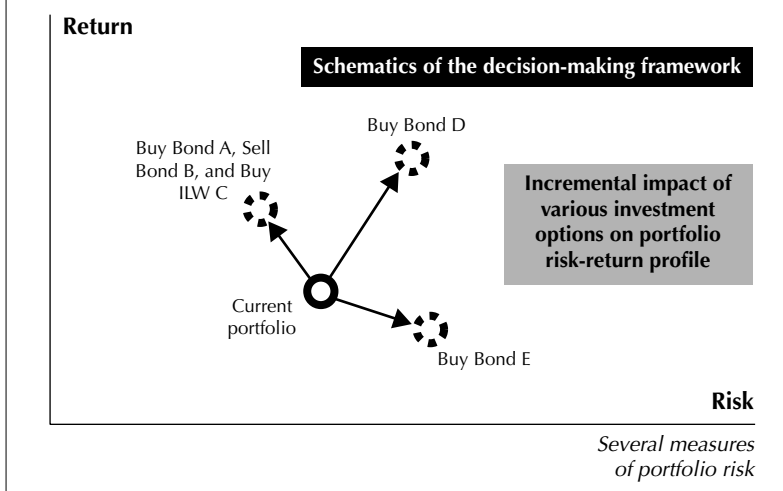


purposes, is not equivalent to that of mean-variance optimisation despite the graphical similarity between Figure 16.6 and the risk–return graphs used in the Markowitz framework. The “Risk” shown in Figure 16.6 is neither variance nor standard deviation. In fact, it is a combination of several risk measures that are presented on the same axis for illustrative purposes only. “Return” too is not necessarily simply the expected return over a one-year horizon, as it can include more than one return measure. The true view of the risk–return trade-offs and the portfolio risk–return profile is a multi-dimensional surface, with the number of dimensions depending on the number of risk and return measures being considered, which is particularly important in the multi-objective optimisation context.

The real decision-making process is considerably more complex but still follows the schematics shown in Figure 16.6 overleaf. Very often, it does end up as a choice between two or three options; and often, the choice is between doing nothing and buying or selling one specific security. These final choices are made after the optimisation software has been utilised, and they involve a significant degree of judgement.

The value of qualitative analysis can be even greater when a very sophisticated optimisation approach with numerous constraints is being utilised. This is exactly the situation where transparency can be absent, and all that is available is the output of the optimiser. The sensitivity to inputs has to be examined very carefully, as well as the sensitivity to slight changes in the

Figure 16.6 Illustrative schematics of the decision-making process based on marginal impact of investment decisions on the portfolio risk–return profile on the multi-dimensional surface



output (security weights in the portfolio). Careful examination of the optimal solution can also be particularly valuable when soft constraints based on local objective functions are used in the optimisation process (see “Portfolio constraints”).

A danger for all optimisers is that they might output an “optimal” decision that might be only a local maximum or minimum of the objective function, while a better solution actually exists elsewhere. There are quantitative techniques to minimise the chance of this happening, but qualitative input can also help to avoid these situations.

Multi-objective optimisation

The multi-objective optimisation process is used when there is a need to simultaneously optimise more than one objective subject to certain constraints, so more than one objective function needs to be maximised or minimised. The objectives are typically conflicting. In the context of portfolio optimisation, we want to maximise return measures while minimising risk measures (again, subject to certain constraints).

The result of the multi-objective process is not a single portfolio allocation, but rather a set of decisions, each of which represents a trade-off among the conflicting objectives. Together, they form what is referred to as the Pareto frontier. This is the set of all solutions that are “Pareto-optimal”; that is, for

each of them an improvement in one objective function can be achieved only at the expense of deterioration in at least one other. Improvement refers to an increase in the objective function that is being maximised or decrease in the function that is being minimised. This is sometimes described in terms of Pareto dominance: a Pareto-optimal solution must be better than another solution in respect to at least one objective function; if it is worse than another solution in respect to all objective functions, it is dominated and is not Pareto-optimal, which means it does not lie on the Pareto frontier.

All points on the Pareto frontier are Pareto-equivalent in the sense that none of them is dominated by others. This is true for the traditional Pareto frontier; but the situation changes if we assign relative weights to the objective functions or introduce some preferences in another way. (See also the discussion of soft constraints above.) The traditional formulation of the multi-objective optimisation problem has all of the objective functions minimised. A problem where some objective functions are minimised (risk measures) while others are maximised (return measures) can be converted to the traditional formulation. We can then see the optimisation problem as obtaining the set of vectors that minimises the set of objective functions $F(F_1, F_2, \dots, F_L)$, where L is the number of objective functions being minimised. Pareto-optimal solutions – those that lie on the Pareto frontier – are also referred to as non-dominated or efficient. The Pareto frontier, in turn, is sometimes called the efficient frontier. It exists in the L -dimensional space, where each dimension corresponds to an objective function. The Pareto frontier by itself does not tell a portfolio manager which portfolio is the optimal one. It represents a set of portfolios that are often infinite, and the portfolio manager has to choose from this set based on criteria not used in the optimisation process that has been utilised to obtain the Pareto frontier.

In a hypothetical example, we might want to maximise the expected portfolio return $E(R_p)$ while minimising the probability of losses (negative returns) $P(R_p \leq 0)$ and the 99% VaR, subject to some constraints and all calculated using a one-year time horizon. The resultant set of solutions, each of which comprises the weights w_i of the portfolio components, is the Pareto frontier for this optimisation process.

If the objective functions and the constraints satisfy the conditions of linear programming optimisation, as outlined in Panel 16.5, then multi-objective linear programming techniques can also be used for solving multi-objective optimisation problems, simplifying the problem. The example above does not fall in this category. The primary issue in trying to apply linear programming to single- and multi-objective optimisation is

typically not the objective functions, as in the example above, but rather the type of constraints used. As mentioned above, TVaR constraints lend themselves to the linear programming optimisation while VaR and many others do not.

There are numerous methods for solving multi-objective constrained optimisation problems. Numerous software packages are available for this purpose. Examples of such methods include the following.

- ❑ Methods combining several objectives in a single aggregate objective function. This allows the use of single-objective optimisers and avoids having to choose a solution from a Pareto frontier since only one solution is generated. Such an aggregate objective function can be a linear combination (sum) of the simple objective functions, with more important objectives having greater weights.
- ❑ Methods that utilise more than one aggregate objective function, such as the successive Pareto-optimisation (SPO) method.
- ❑ Evolutionary algorithms (EA). Multi-objective optimisation evolutionary algorithms (MOEA) are very flexible and offer promising approaches to portfolio optimisation. Unfortunately, computational complexity prevents their wide use. The MOEA family includes quite a number of algorithms. Genetic algorithm (GA) is the most popular.¹³ Particle swarm optimisation (PSO) is another that is often used.
- ❑ Simulated annealing (SA) methods.¹⁴ SA are very general stochastic methods; unfortunately, the specific algorithms used are not very efficient. These methods have certain advantages in dealing with discrete variables, which are sometimes found in the case of cat insurance risk.

Other methods exist as well, but their use in portfolio optimisation is problematic at best.

PITFALLS OF STANDARD OPTIMISATION TECHNIQUES

Men get into trouble by taking their visions and hallucinations too seriously.

H. L. Mencken

Extreme caution is needed when using any optimisation method or software unless they are well understood and any potential weaknesses are clear. Optimisation techniques, when used improperly or uncritically, can produce solutions that are far from optimal and portfolios with little return or a lot of risk.

First of all, we need to examine the sensitivity of the optimal solution to optimiser inputs. Whole classes of optimisers have been called error

maximisers, as they amplify input errors. Having a low degree of certainty that the loss distribution for a particular cat insurance risk is correct has to be taken into account in optimisation. Quite often, an optimiser produces “optimal” portfolios that are dominated by a handful of securities (and sometimes only one security). A likely explanation then is that the risk of these securities is understated and/or the return is overstated.

An optimiser may produce an “optimal” portfolio that does maximise (or minimise) the objective function, but the maximum (or minimum) is local, and a better solution is missed altogether. This is a problem with most optimisers, but there are steps to minimise the chance of it happening.

Some of the problems mentioned above are not unique to managing portfolios of insurance-linked securities but have plagued many a manager in other asset classes. A few issues are specific to ILS. First, the use of proper optimisation techniques is not common. Second, there is a greater probability of mistakenly choosing a portfolio corresponding to a local maximum (or minimum) of the objective function. Finally, the non-modelled risk can be another potential issue.

REMODELLING AND PORTFOLIO OPTIMISATION

The three modelling companies – AIR, EQECAT and RMS – provide risk analysis of outstanding cat bonds along with their software packages for ILS investors. RMS also provides analysis of extreme mortality bonds, which can be mapped in its software along with property cat bonds.

The modelling firms makes information available on all outstanding cat bonds (with some exceptions for one of the three firms). When a new cat bond is being issued, all three firms will likely have their risk analysis available before the bond is priced.

The information on the outstanding cat bonds – the so-called remodelling data – is the result of each of the firms having examined offering documents and making an attempt to produce its own risk analysis based only on information available to investors, as opposed to the more extensive data used in the risk analysis included in investor documentation for the bonds.

The remodelling serves three very important functions that are essential for portfolio management.

- Remodelling data provides investors with the ability to map all the securities together and obtain a loss distribution for a portfolio without having to make any correlation assumptions.

- ❑ Remodelling data maps exposure for each bond. This makes the analysis much easier for parametric bonds and is even more important for indemnity bonds, where there is a need to make some assumptions to determine exposure distribution at a more detailed level than shown in the OC. Most investors cannot do this on their own; and for others, it provides an important reference point to which they can compare their own exposure mapping.
- ❑ Remodelling data allows us to bring the same common denominator to the risk analysis for all securities, as if all the bonds had been modelled for the OC by the same modelling firm.

Even the modelling firm that has provided official (included in the OC) risk analysis for a specific indemnity bond, and has had access to very detailed exposure information, uses only the summarised data from the OC in its remodelling analysis. (The more cynical investors, however, have questioned the truth of this statement.)

Investors still need to make adjustments to the risk mapping in the risk analysis when they have (or think they have) superior information. Or an investor might have views on biases of some models, and would like to make adjustments accordingly. Remodelling addresses this last scenario only partially, since an investor would usually have views not on model biases in general but rather on specific peril/territory combinations. Remodelling thus does not eliminate the need for some adjustments for the investors who believe they have an informational advantage.

While the main advantage of remodelling data that comes with the software is the improved ability to model ILS on a portfolio basis, the commentary provided in the reanalysis is often insightful and can add significant value in the analysis of individual securities.

SENSITIVITY ANALYSIS AND SCENARIO TESTING

For every complex problem there is an answer that is clear, simple, and wrong.

H. L. Mencken

A portfolio optimiser might take all the required inputs, employ a sophisticated optimisation algorithm for the calculations, and come up with an “optimal” portfolio. This portfolio might be truly optimal for the investor. Or it might not be.

Some of the potential pitfalls of the portfolio optimisation process have been mentioned above. For example, some categories of optimisers can easily become error maximisers, amplifying errors in the input data and

producing portfolio decisions that are far from optimal. Sensitivity analysis can reduce the probability and magnitude of such errors.

As in managing portfolios of most securities, one of the goals of sensitivity analysis is the development of a systematic approach to changing input parameters to determine which of them produce disproportionate impact on the model output. For cat ILS, this process can start before the optimisation, as we can perform sensitivity analysis of the current portfolio using a model such as CATRADER, eCAT or Miu. It is often immediately clear which inputs might have a disproportionate effect on the portfolio loss distribution. These would tend to be the ones linked to the peak perils; or the ones where there is a significant chance of mispricing (such as when the investor has a very different view of the risk than the one presented in the OC). A disproportional effect would also occur when the model is not sufficiently sensitive to changes in one of the parameters.

Sensitivity analysis in examining the optimiser output has the same parameters to consider, but it is also important to be aware of the specific optimisation algorithms used and their potential to introduce certain biases or produce certain types of solutions that lack robustness.

Scenario testing is also very important in portfolio management and the analysis of the solutions produced by an optimiser. It can help to recognise mistakes, determine sensitivity to some parameters, discover portfolio behaviour under unusual conditions, identify unexpected correlations and dependences and provide an additional reality check for the portfolio-management and optimisation process. The two main types of scenario testing useful in this context are the following:

- ❑ analysis of realistic scenarios and their impact on the portfolio; and
- ❑ stress testing to understand portfolio behaviour in extreme circumstances.

The scenario-testing framework for catastrophe ILS is the same as that used for portfolios of other securities. It is the scenarios themselves that are different, as they can include, for example, the impact on the portfolio of a particularly devastating natural catastrophe.

ADDITIONAL CONSIDERATIONS

There are many other considerations involved in the management of portfolios of catastrophe insurance risk. Identification of the “hidden” risks is one of them. Evaluation of the dependence of ILS pricing on the capacity available in the reinsurance markets is another. Additional considerations include the following.

- ❑ Many catastrophe insurance risks exhibit a clear seasonality effect that is also present in the pricing of insurance-linked securities, and that is relevant to the construction and management of an ILS portfolio.
- ❑ Simple approaches to risk and portfolio management, though they are often considered “naïve”, can still be useful as a check on the more sophisticated methods for managing portfolio risk. For example, the “risk bucket” approach, where portfolio risk is decomposed into combinations of peril and territory, can identify mistakes in modelling.
- ❑ Besides cat insurance risk, catastrophe ILS can be exposed to a number of other risks, ranging from market to credit. These risks are very difficult to quantify and, in some cases, even to identify. Still, they have to be taken into account in the analysis and portfolio decisions.
- ❑ It is important to have a view on future developments in the cat ILS market, as they can affect supply of new issues and change pricing levels. The reinsurance market and its level of capacity for cat risk can have an important impact here. So can financial market crises, but these are more difficult to foresee.
- ❑ Liquidity considerations and risks having to do with liquidity can be easily overlooked, but they too can be important; liquidity risks have to be properly assessed and liquidity management policies established and followed. This consideration is particularly important for investors in securities where collateral requirements can suddenly change (such as margins for exchange-traded cat derivatives) or funds that can face significant redemptions.
- ❑ Monitoring developments in the cat modelling world can help identify certain modelling biases and allow the astute manager to be one of the first to become aware of upcoming changes in the models and their effect on the assessment of specific risks. This can be a significant source of competitive advantage.

Many other considerations are important in cat ILS portfolio management. The list is long. The additional numerous considerations are significant and should not be neglected. At the same time, they should not take the focus away from the key performance drivers of a cat ILS portfolio.

PERFORMANCE MEASUREMENT

However beautiful the strategy, you should occasionally look at the results.

Winston Churchill

Performance measurement is very difficult when analysing a manager or fund. Often the only information available to an investor, besides some qual-

itative description, consists of the monthly or quarterly returns and the standard performance measures based on them. These standard measures are often inapplicable to the analysis of a portfolio of catastrophe risk. They are, however, very important to the managers, since many investors will still be basing their decisions on these standard measures.

Chapter 17 has a detailed description of the key performance statistics used in the analysis of hedge fund investment results. These measures include the following:

- ☐ average return;
- ☐ compound monthly return;
- ☐ compound annualised return;
- ☐ active premium;
- ☐ monthly standard deviation;
- ☐ annualised standard deviation;
- ☐ downside deviation;
- ☐ longest drawdown;
- ☐ maximum drawdown;
- ☐ monthly Sharpe ratio;
- ☐ annualised Sharpe ratio;
- ☐ Sortino ratio;
- ☐ Treynor ratio;
- ☐ Calmar ratio;
- ☐ Jensen's alpha; and
- ☐ gain-to-loss ratio.

While there is a strong argument that most of these performance measures are not particularly meaningful for cat ILS, the reality is that many investors will still use them, so these measures cannot be neglected.

Track record can be misleading

If a manager's strategy includes delivering high risk-adjusted returns with low correlation to the rest of the financial markets, it is easy to notice the cases when the correlation is still high. The high beta results for the near-zero beta promise are not hard to spot. The difficulty arises when the returns are relatively high and the volatility is low. Is this a sign of a good fund manager? The answer could be negative or inconclusive.

In managing catastrophe risk, the above-described "time-bomb" events pose the real danger. Relatively high absolute returns and low volatility do

not mean that the returns are high on the risk-adjusted basis. In fact, abnormally high absolute returns may be a sign of poor risk management. The manager can be taking very significant risk without reporting it to investors or even realising the true risk exposure of the portfolio.

The track record can appear perfect “until it doesn’t”, when a hurricane or earthquake wipes out a sizable portion of the portfolio. The claims of the possible loss from a single catastrophic event being limited to a certain percentage of the portfolio are difficult to verify. The portfolio might contain risks that are not immediately obvious, and these risks can remain dormant for a number of years. Even claims that the portfolio “went through the hurricane Katrina season without any losses” have very little value. The nature of the beast – the catastrophe – is such that the events testing the portfolio and risk management do not come often, and a manager can maintain a seemingly perfect track record while doing almost everything wrong.

The problem is even bigger on the life insurance side. For example, portfolios of life settlements that do not carry catastrophic risk in the traditional sense can still lead (and have led) to investor losses of a catastrophic nature. Without a liquid market to ascertain proper value of these investors – or at least to serve as a reference point – mark-to-model is the approach used in valuing these assets. In this corner of the ILS market, full of naïve investors and naïve managers, it is possible for managers to lead themselves and others into thinking the portfolio is doing very well, when in reality losses are mounting. It could take many years for some of these managers to realise or to admit to their investors that significant write-downs are necessary. Until that happens, they might continue to collect incentive fees. The mark-to-market approach, out of necessity replaced by mark-to-model, becomes mark-to-make-believe.

CONCLUSION

Managing catastrophe risks on a portfolio basis is as important as proper analysis of individual securities. In fact, sometimes it can add more value. By the same token, it is clear that when it is not performed properly it can lead to significant losses. Avoiding significant losses, however, is only one of the goals of portfolio management; in fact it is more properly considered as a constraint to be followed, than as a goal. The goal is generating high risk-adjusted returns based on the investor risk appetite and return preferences.

Risk management is an essential element of portfolio management. In fact, it can be accomplished only when looking at the whole portfolio. Risk-adjusted return is relevant to the investors only in the portfolio context; it can be maximised, and even measured, only for the whole portfolio.

Managing cat ILS on a portfolio basis is still an emerging field, and many investors have neither the necessary tools nor the expertise. This situation is likely to change rapidly as the process of investor education continues and new modelling tools become available. Cat modelling software designed specifically for cat ILS investors is improving and new features are being constantly added. The powerful portfolio optimisation tools available today are used by only a few ILS investors, but this too is likely to change. The level of expertise in the analysis of individual cat ILS, and in their management on a portfolio basis, is growing, albeit relatively slowly.

Recognition of the importance of the true portfolio approach is growing, as well. It is becoming more widely recognised that the ability to effectively manage a portfolio of cat risks can be a valuable source of competitive advantage and an important differentiator in this still highly inefficient marketplace.

- 1 There are some rare exceptions to this statement. Under MPT, the statement is always true.
- 2 There is considerable confusion in the terminology when it comes to the concepts of alternative and exotic beta. See "Beyond active alpha" by Bob Litterman (Goldman Sachs) for the definitions and explanations with which this representation is consistent. Another form of exotic beta, that of applying "exotic" strategies to traditional asset classes, is not mentioned since it is not relevant to this discussion.
- 3 In an occurrence insurance policy, second-event coverage refers to the indemnification of losses for a second occurrence of a qualified event. Similarly, in ILS second-event securities assume the risk of the second qualified event occurring during the coverage period. In this case, only the second event is covered; losses from the first event are not reimbursed. See Chapter 3 and Chapter 5 for additional discussion of this topic.
- 4 While some defined all these measures of risk exactly the same way and use the terms tail value-at-risk (TVaR), conditional value-at-risk (CVaR) and conditional tail expectation (CTE) interchangeably, the original definitions of the terms differed between TVaR and CVaR on the one side and CTE on the other. In some literature, there is also a distinction made between the definitions of TVaR and CVaR, but the vast majority of practitioners and academics do not differentiate between the last two terms.
- 5 Taking significant risk even for a short period of time will of course be reflected in the VaR and TVaR calculated over a much longer time period.
- 6 The rating process is complex and involves qualitative inputs. The models maintained by rating agencies are the main but not the only determinant of the ratings. Without engaging in a dialogue with rating agencies, it is not always possible for a company to determine the precise amount of catastrophe risk that it can take.
- 7 TVaR is easiest to interpret in the regulatory context. For example, when this concept is used in the context of solvency management for a (re)insurance company (as opposed to considering only its cat risk portfolio), it is very close to the concept of expected policyholder deficit, which is the expected loss beyond the probability of ruin that regulators would have to deal with through the system of guarantee funds, or in other ways if it occurs.
- 8 The reason for such constraints is the market inefficiency that manifests itself in the broker quotes being only "indicative" and the limits on how much of a particular security an investor can buy at a certain price or at all. While liquidity in the cat bond secondary market continues to improve, the market cannot be called liquid at this point.

- 9 Qualified institutional buyer status.
- 10 There is no standard compensation structure for ILS funds as there is no standard structure in the alternative asset management industry in general. The 2-and-20 compensation arrangement has been under attack for a long time, and most funds are accepting the reality of investors not willing to pay fees that high. However, it is likely that many hedge funds will continue to be able to charge fees based on the 2-and-20 structure. Dedicated ILS funds have found it very difficult to justify the 2-and-20 structure to investors.
- 11 For only catastrophe bonds, the relatively high probability of achieving returns greater than 20% and 30%, as shown, is unlikely, unless the existing holdings in the initial portfolio were bought at an opportune time. The overall relationship between risk and reward also appears to indicate the presence of other types of ILS. The degree of smoothness of the distribution function (see also Figure 16.5) is another reason to suspect that more than cat bonds are included in the portfolio.
- 12 In reality, the changes are incremental rather than marginal since most of the changes can be implemented only in increments rather than be continuous functions. For example, there are minimal limits on how much of an individual security can be bought or sold; these limits are typically higher for ILS than for traditional securities.
- 13 L. Zeng was probably the first to report the use of GA in this context.
- 14 In some ways, SA methods can very closely resemble evolutionary algorithms such as PSO.

Managing Portfolios of Multiple Types of ILS

This chapter explores the topic of managing portfolios of various types of insurance-linked securities, where only some of the insurance risk is of the catastrophe type, or where ILS instruments that are designed to have very low correlation with financial markets are combined with securities that have significant risks in addition to the pure insurance risks.

Combining insurance-linked securities of very different types – such as cat bonds and embedded-value securities – is not common, but some dedicated ILS funds have pursued this approach. This chapter discusses the issues that arise in managing such ILS portfolios – in particular, issues that are not present in managing portfolios of only catastrophe insurance risk (examined in Chapter 16).

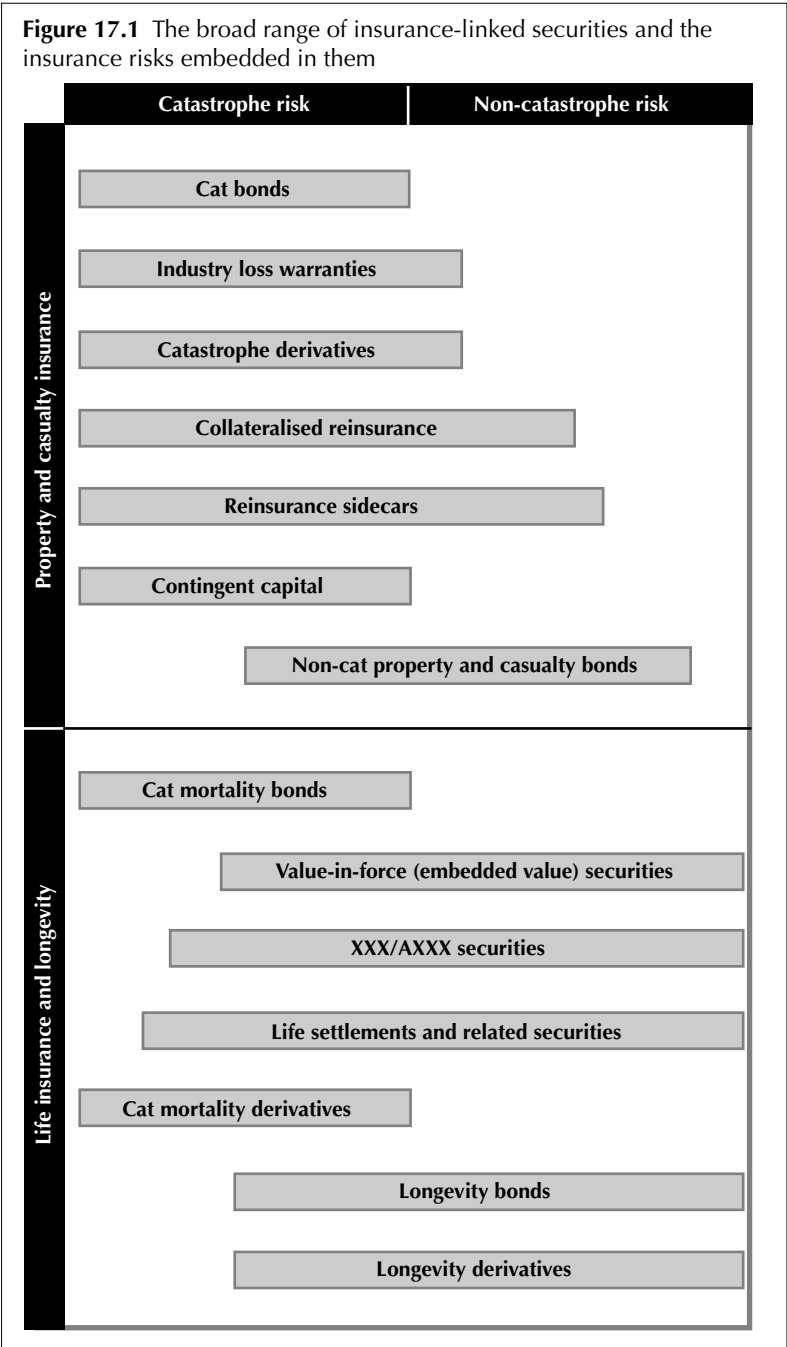
TYPES OF INSURANCE-LINKED SECURITIES

While the universe of ILS instruments is very broad, the ILS markets tend to be segmented. Some investors specialise in cat bonds only; some have interest in other specific cat ILS or in the whole cat ILS market; some deal only in life settlements; and there are others who have only invested in Regulation XXX securities. Only a handful of investors in property cat ILS also invest in mortality- or longevity-linked securities other than extreme mortality bonds.

Figure 17.1 repeats the classification of main types of insurance-linked securities that was used in Chapter 2. It shows that there is rarely a clear border between catastrophe and non-catastrophe insurance risk. Many types of ILS can cover most of the spectrum between the point of exposure to the far tail of the insurance loss distribution, and the point where the insurance risk involved is clearly not of the catastrophe type.

The list in Figure 17.1 is far from exhaustive, and the chapters devoted to individual types of insurance-linked securities have included even more

Figure 17.1 The broad range of insurance-linked securities and the insurance risks embedded in them



varieties. In addition, the statement that many types of ILS cover a large part of the spectrum between cat and non-cat risk is true in two ways:

- ❑ some securities can transfer cat risk to the capital markets while others, in the same general category, transfer non-cat risk; and
- ❑ there are securities that have both cat and non-cat insurance risk components embedded in them.

RATIONALE FOR COMBINING DIFFERENT TYPES OF ILS IN THE SAME PORTFOLIO

Not everybody agrees that it makes sense to combine multiple types of ILS in the same portfolio. In examining the pros and cons of this strategy, it is also recognised that what may make sense for one investor may not be reasonable or even possible for another, all for very objective reasons. The discussion is focused not on combining similar ILS types, which is generally seen as highly beneficial, but more on putting together and managing multiple types of ILS securities that have rather different characteristics.

Flexibility and diversification

Diversification within the ILS asset class is sometimes brought up as a reason for combining different types of ILS in the same portfolio. Adding new types of risk has the effect of lowering the overall risk of the portfolio, which is the standard diversification argument in portfolio management. Obtaining exposure to the risk of longevity should have a diversifying effect on a portfolio of catastrophe risk and improve its risk-adjusted return.

In addition, having access to a broader universe of investment instruments gives the portfolio manager extra flexibility in security selection and the overall portfolio optimisation. It also allows the manager to increase or reduce allocations to individual ILS classes to take advantage of the changes in market conditions.

Skill set and expertise

It is important to consider the availability of expertise needed to analyse different types of insurance-linked securities. For example, property catastrophe insurance risk is very different from the risk of longevity, and a property cat modeller is unlikely to have any knowledge of, let alone expertise in, the risks of longevity. On the other hand, if expertise in various types of insurance-linked securities is available, it may become a significant source of competitive advantage (assuming that combining multiple types of ILS makes general sense for the specific investor).

The issue of expertise is even more important in portfolio management of different types of ILS; this skill set is different from the skill sets needed to analyse individual securities of one type or another or of managing portfolios of only one type of ILS.

CORRELATION AMONG DIFFERENT TYPES OF ILS

It is far better to foresee even without certainty than not to foresee at all.

Henri Poincare

Correlation among different types of ILS is not always easy to assess. Dependence issues of catastrophe insurance-linked securities have been examined in the previous chapter and in Chapters 3 and 5. The risk of longevity or non-cat mortality appears to be uncorrelated with the risk of natural catastrophes.¹ Some correlations can be quite noticeable, for example among some embedded value securities and those that are linked to excess reserves or longevity. These have to be correctly taken into account in the modelling process to properly reflect the risks of the overall portfolio.

Many non-cat insurance-linked securities have a rather strong market risk embedded in them along with the insurance risk. Through their exposure to market risk these ILS are also correlated with each other. The correlation with the markets also has to be properly reflected in the models; constraints on the risk measures associated with it might be more or less restrictive in some cases, depending on the investment objective.

TENOR AND LIQUIDITY

The issue of security tenor and duration is one of the thorniest in combining some types of insurance-linked securities. It is not even always related to the interest-rate risk present in most if not all longer-term securities. The idea of combining Regulation XXX securities maturing in 25 years, with cat derivatives positions that are expected to be settled in seven months, creates difficulties in portfolio management when the two are managed as part of the same portfolio; and in the minds of most investors this defeats any possible advantages of such combination. The issue becomes even greater when life settlements or similar securities are added to the mix. Life settlements have significant interest rate exposure and their maturity is not always easy to estimate. Embedded value has some of the same problems.

In general, there is nothing wrong with investing in both long- and short-duration instruments. The problem comes in managing them in the same portfolio as if they were similar instruments, since this does create serious difficulties in portfolio optimisation. Given the very low liquidity of most of

the long-duration insurance-linked securities, the time horizon used in portfolio optimisation changes significantly – a number of time horizons have to be added and often reinvestment assumptions have to be made. Few can even enter such a game, since investing in illiquid securities with very long duration is not a good choice for a fund with no lockup provisions or with lockups limited to only two or three years. On the other hand, having part of the portfolio invested in shorter-term, more liquid ILS can provide the flexibility to investors whose portfolios are invested primarily in illiquid instruments such as life settlements. This could be seen as another potential advantage of combining more than one type of ILS in the same portfolio for such investors: liquidity concerns are reduced through investing in another type of ILS instead of other securities or cash, maintaining a high level of the exotic beta attraction of the overall investment.

Liquidity-management issues gain significant importance in such cases, in particular when negative cashflows are expected in the first years, as is the case for life settlements. Strict liquidity-management policies and careful monitoring of the portfolio are required then to avoid liquidity-related problems.

PORTFOLIO OPTIMISATION

The problem of portfolio optimisation is not unique to portfolio management of multiple types of insurance-linked securities. The constraints, however, can differ, often significantly, to reflect the risks involved. The objective function can differ as well, but to a lesser degree.

A common approach – due to its simplicity – is to have broad asset allocation established first, based on either some type of optimisation or qualitative considerations (such as liquidity constraints, investment time horizon, expectations of future developments in some ILS markets, and the level of in-house expertise in various ILS areas); then each of the subportfolios is optimised separately, based on the constraints and goals established for each one (with each subportfolio dedicated to a distinct type of ILS). The results of the optimisation at the subportfolio level can then be used to reoptimise the total portfolio. This process can go through several iterations.

An approach where the optimisation is performed at the portfolio level, taking into account all types of ILS in the same optimisation process, may be more correct from the theoretical point of view but is very difficult to implement in practice. The sensitivity to inputs, and even more so to constraints, becomes very high, while the choice of constraints is far from being intuitive. The resulting solution would often exclude most ILS types from the optimal portfolio. This might or might not be correct.

The challenges of optimising portfolio of ILS of multiple classes

In managing portfolios including several types of insurance-linked securities, the challenges of creating a true optimisation engine to reflect the risk constraints and decision variables can be daunting. Unlike the case of optimising a portfolio of catastrophe ILS, here there is no neat map of exposure and results for every scenario in the simulation from the standard model.

A particular difficulty is the wide disparity between the tenors of the securities in the same portfolio. To have some securities maturing in a year and others in 25 years is unusual, and they are difficult to model together.

Correlation among securities of various types is another thorny issue that cannot be resolved as cleanly as when limiting the analysis only to cat ILS (as described in Chapter 16).

The number of parameters and variables used in the optimisation process grows as well, which can make traditional optimisation computationally impossible; then there is a need for simplifying assumptions or new approaches to portfolio optimisation.

Sensitivity analysis

When analysing and optimising the total portfolio, sensitivity analysis also plays a greater role than it does when focusing on only one ILS type. Portfolio effects can be very unusual.

Of particular importance in this type of sensitivity analysis is the examination of how sensitive the solutions are to the choice of constraints used in the optimisation process, which in such cases is usually much greater than the sensitivity to actual inputs. The choice of the types of constraints (variables) and their values can drive the optimisation process, with small changes resulting in dramatically different optimised portfolios. A small change in one constraint, or a substitution of it with a seemingly similar constraint, can result, for example, in a change in asset allocation from 75% in one type of ILS to this type being completely excluded and the balance being shifted to other insurance-linked securities. The more intuitive the constraints, the better. This complicates the use of tail value-at-risk (TVaR) since choosing its specific values as a constraint is far from intuitive and requires considerable judgement, but the sensitivity to this choice can be great. The lack of robustness can be a persistent problem.

THE ARGUMENT AGAINST COMBINING ILS OF MULTIPLE TYPES IN THE SAME PORTFOLIO

The discussion above reveals numerous issues and potential pitfalls in investing in insurance-linked securities of different types. The three most important of these are the following:

- ❑ problems with combining short-term with very long-term securities;
- ❑ introducing greater correlation with the global financial markets that can be limited by assembling insurance-linked securities only of specific types; and
- ❑ greater expertise needed for the analysis of ILS of different types and their management as one portfolio.

Assuming an investor is even allowed to invest in long-term and relatively illiquid securities, the most important difficulty is the third factor, as combining multiple types of ILS significantly increases the overall complexity of the management task and requires a skill set that is rarely found.

All of this does not mean that different types of ILS should not be combined and managed as one portfolio. The argument goes against only certain combinations, in particular those where short-term and very long-term securities are combined, and applies especially when the analytical skill set is not present.

PORTFOLIO VALUATION ISSUES

In a portfolio of insurance-linked securities the portfolio valuation issues differ by the type of ILS or subportfolio. As with other securities, the valuation should be based on market prices when such are available and can be considered reliable. For some insurance-linked securities, this is the case. For example, for catastrophe bonds, such prices are available from several dealers for almost every single security. The prices are not firm but rather indicative; nonetheless, they provide a good market input into valuation. (The question of whether to use the simple average of pricing indications from several sources, or a different measure, is not unique to ILS: it is common to all securities that trade rarely and lack readily available firm prices.) For exchange-traded securities, such as IFEX cat derivatives, exchange-reported settlement prices are directly used for valuation purposes even though they often do not represent the actual prices of the last transaction.²

For many types of ILS, there are no readily ascertainable market prices from reliable sources, and inputs from the models have to be used for valuation purposes, either exclusively or to supplement pricing data on these or related securities. The approach is exactly the same as for other non-liquid securities; the difference comes from the types of models being used and the inputs in these models. For mortality-linked securities linked to a specific group of insured individuals, these inputs (the LEs in the broader sense of

the total sets of applicable mortality rates) are of great importance and drive the valuation. Such inputs have to be periodically reassessed as opposed to simply using in valuation the assumptions that were used when the securities were initially purchased.

Valuing equity investments in reinsurance sidecars presents obvious difficulties. Sometimes these are addressed, if enough information is available to investors; but often they are valued at the original cost.³

Discount rate used in the valuation of ILS that have expected cashflows spread over a long time period is linked very directly to the general financial markets, and it can have quite a significant impact on the valuation. For insurance-linked securities that have expected negative cashflows in the first years (something not often seen in other asset classes), this issue is also linked to the effective cost of having access to the extra funds in the future, introducing additional complexity.

Cat bonds largely avoid the problem of interest-rate risk sensitivity, as they are almost always floating-rate instruments; but this concern is not important in valuation since market prices are usually available for these securities.

PERFORMANCE MEASUREMENT

The comments on the inapplicability of many investment performance measures to cat risk portfolios (Chapter 16) are equally relevant for investment portfolios that include other types of insurance-linked securities, though for different reasons.

The track record of a cat ILS portfolio manager can be misleading simply because in the fund's short history there have not been catastrophic events to cause any sizable losses. This problem can rarely be found in other asset classes, at least to the same extent. In fact, managers with the best records are often the ones who take on unreasonable risk. The records can be misleading, but they are true and almost always can be easily verified.

For many non-cat ILS, there are no sufficient market inputs to implement proper mark-to-market valuation, and mark-to-model is used instead. In addition, these securities usually have long (and often quite uncertain) duration. If the modelling assumptions are not re-examined, the reported results for these securities will continue to mirror, relatively closely, the initial expectations based on the same model. The track record might appear rather attractive, but it will not be true if the valuation has not been done properly (that is, LEs not adjusted when necessary). In this case the reported results will be wrong, distorting the track record. Unlike the case of cat ILS, this type of distortion is not unique to insurance-linked securities.

The argument of inapplicability of standard performance measures to cat ILS is weaker for non-cat ILS, especially when mark-to-market valuation methods can be used. Many non-cat ILS have significant levels of traditional market risks embedded in them, leading to market-linked volatility. In the analysis of portfolios of multiple types of insurance-linked securities, standard performance measures – the ones that are routinely used by many investors in funds – are more meaningful than for cat-only insurance-linked securities. Some of these measures, only listed in the previous chapter, are defined below.

Some standard performance measures

The main measures of portfolio return used in hedge fund performance analysis are defined in Panel 17.1.

Next, Panel 17.2 describes some of the measures of portfolio risk derived from historical return. Some of these measures can be defined in slightly different ways. For example, downside deviation has a number of definitions, all producing rather similar results. Sometimes downside deviation is considered to be equal to semi-variance, which is usually defined differently. Annualisation can be done in more than one way.

The R_{MAR} referenced in the formula for downside deviation, and later in the definitions of risk-adjusted measures of return in Panel 17.3, is the minimal accepted return level. It is often assumed to be equal to zero, or to the risk-free rate R_f .

Here too the same measure can often be defined in more than one way. Regardless of what appears to be the most “theoretically correct”, the definitions most often used in practice in hedge fund reporting are the ones shown in Panel 17.3. Often, there is no consensus even in the way this practical reporting is done. For example, while most seem to use the mean return in their Sharpe ratio calculation, others use the geometric (compounded) average.

Sortino ratio is the analogue of the Sharpe ratio that uses downside deviation instead of the standard deviation of returns, thus focusing on the downside risk (bad volatility) without imposing a penalty for positive deviations from the mean (good volatility).

In the analysis of investment performance, we need to keep in mind that many of these measures reflect the risk of catastrophic events only to a small degree, and for non-cat mortality-linked and other securities even key risks can become evident only after a long period of time. This again confirms that the track record can be very misleading in the performance analysis of an ILS fund.

PANEL 17.1 SOME STANDARD MEASURES OF PORTFOLIO RETURN

For a regular investment portfolio the measures of risk and return are typically calculated based on the historical performance. Thus, the data is discrete. The formulas below show how to use this discrete data for calculating appropriate risk and return measures.

In the definitions below, the discrete data is monthly, based on the assumption that the fund strikes its net asset values (NAVs) on a monthly basis. Thus, R_k is the return for month k , while n is the total number of months being considered.

$$\text{Average return} = \frac{\sum_{k=1}^n R_k}{n}$$

$$\text{Compound monthly return} = \sqrt[n]{\prod_{k=1}^n (1 + R_k)} - 1$$

$$\text{Compound annualised return} = (1 + \text{Compound monthly return})^{12} - 1$$

$$\begin{aligned} \text{Active premium} &= \text{Portfolio annualised return} \\ &\quad - \text{Benchmark annualised return} \end{aligned}$$

PANEL 17.2 SOME STANDARD MEASURES OF PORTFOLIO RISK

For a regular investment portfolio the measures of risk are typically calculated based on

$$\text{Monthly standard deviation} = \sqrt{\frac{\sum_{k=1}^n (R_k - \text{Average return})^2}{n - 1}}$$

$$\text{Annualised standard deviation} = \sqrt{12} \text{ Monthly standard deviation}$$

$$\text{Downside deviation}_{MAR} = \sqrt{\frac{\sum_{k=1}^n (R_k - R_{MAR})^2}{n}}$$

$$\text{Longest drawdown} = \text{Max}_{\text{all drawdowns}} (\text{Drawdown length}_m)$$

$$\text{Maximum drawdown} = \text{Max}_{\text{all drawdowns}} (\text{Percentage drawdown}_m)$$

Despite how strong the argument is for not using many of these measures in the performance analysis of portfolios of insurance-linked securities, the fact remains that many investors have poor familiarity with this asset class

PANEL 17.3 SOME STANDARD MEASURES OF PORTFOLIO RISK-ADJUSTED RETURN

For a regular investment portfolio the measures of risk and return are typically calculated based on

$$\text{Monthly Sharpe ratio} = \frac{\text{Average return} - R_f}{\text{Standard deviation}}$$

$$\text{Annualised Sharpe ratio} = \sqrt{12} \text{ Monthly Sharpe ratio}$$

$$\text{Sortino ratio}_{\text{MAR}} = \frac{\text{Compound monthly return} - R_{\text{MAR}}}{\text{Standard deviation}_{\text{MAR}}}$$

$$\text{Treynor ratio} = \frac{\text{Average return} - R_f}{\beta}$$

$$\text{Calmar ratio} = \frac{\text{Compound annualised return}}{\text{Maximum drawdown}}$$

$$\text{Jensen's alpha} = (\text{Average return} - R_f) - \beta (\text{Average benchmark return} - R_f)$$

$$\text{Gain-to-loss ratio} = - \frac{\text{Average gain in gain period}}{\text{Average loss in loss period}}$$

and are likely to base their decisions on the traditional performance statistics. Consequently, these measures cannot be completely neglected even by a fund manager who considers them to be completely irrelevant.

INVESTMENT MANAGEMENT POLICY

Consideration of structures and formalised policies in investment management is even more important – and significantly more complex – in the management of portfolios of insurance-linked securities of multiple types than it is in the portfolio management of only catastrophe risk. The infrastructure required for effective investment management likewise is more complex.

While formal policies are expected in investment management in general, regardless of asset class, in this situation simply having weekly meetings of the investment committee is not sufficient: instead a comprehensive framework has to be put in place. The process must include the function of risk management discussed below; it also must ensure that investment decisions are linked with risk management and measurement, and that all parts of the portfolio are being considered in the decision-making process.

In the case where there are managers responsible for subportfolios, each of which comprises one main type of insurance-linked securities, it is essential to ensure that the decisions are made on an integrated basis, taking into account all parts of the overall ILS portfolio and their interaction.

The risks – of losses and of missed opportunities – are so much greater than in the case of cat-only ILS portfolios that a highly formal process has to be put into place and closely followed. This somewhat constrains the flexibility of the portfolio management decision-making process.

The specific formal policies depend on the investor and the risks involved; the elements that have to be present, in addition to those standard to the asset-management industry, can differ significantly from fund to fund, depending on the investment goals, strategy, legal and other constraints, the types of ILS in the portfolio, the level of expertise available and the general infrastructure already in place.

RISK MANAGEMENT

Risk comes from not knowing what you're doing.
Warren Buffett

Risk management in investment portfolio management of insurance-linked securities of multiple types includes all the elements of a traditional risk management framework, along with the additional elements having to do with catastrophe insurance risk (described in Chapter 16). It should also include the elements reflecting the specific risk arising from combining multiple types of ILS in the same portfolio, as described above. As an example, liquidity- and cash-management policies might be needed to address the specific risks of the life settlement subportfolio becoming a greater-than-expected cash drain on the rest of the portfolio.

As in the case of portfolios of cat ILS, risk management, in addition to serving the risk-control function, has to be part of the overall portfolio-management process aimed at maximising risk-adjusted returns. Measuring risks in this case is more complicated and the risk matrices greater, as a number of very diverse risks have to be captured. Most of the non-catastrophe ILS have significant exposure to various market risks in addition to the “pure” insurance risk. The correlations in the portfolio are stronger and can come from more than one source. This is particularly true of the tail events (not only for insurance risk), where the correlations become much stronger.

As additional types of ILS are added to the portfolio, the types of hedging tools and techniques need to expand as well. All of them – from buying an

ILW to entering into a longevity derivative transaction – have to be considered both in risk management and making buy-and-sell investment decisions, as the two are truly interdependent in an expertly managed portfolio.

Stress testing and portfolio monitoring

In the types of portfolios being discussed, correlations and dependencies are often difficult to properly reflect in modelling, and parameter uncertainty in general is much greater. Stress testing, along with sensitivity analysis, then becomes an even more important part of risk management.

It is important to construct stress scenarios that properly reflect the risks of a diversified ILS portfolio. For example, a systemic bias in LEs for the life-settlement part of the portfolio can have an unexpected effect on the other parts of the portfolio, creating a cash strain and a possible need to sell securities of other types while at the same time restricting the ability to rebalance the portfolio.

Portfolio monitoring in managing cat ILS is a relatively straightforward process. When other types of ILS are added to the mix, the overall complexity grows. Small changes in the portfolio can indicate broader developments, which is rarely the case in cat-only insurance-linked securities. In addition, such changes can also help to reveal interdependencies contributing to the overall portfolio risk.

CONCLUSION

Investing in insurance-linked securities is a very specialised field, with further specialisation within it. The reasons for this are historical and have to do with how these markets have developed. They also have to do with differences in the skill set needed for analysing different types of insurance-linked securities and managing their portfolios. Finally, for most investors, combining very different types of ILS in the same portfolio simply does not make sense: it adds questionable if any value but introduces new risks. For some investors, only certain types of ILS can be used to achieve their investment objectives; there might even be legal restrictions on investing in some types of insurance-linked securities.

Particularly questionable are the so-called ILS “fusion” strategies, where combining insurance-linked securities (for example, life settlements) with investments such as project finance or distressed debt is purported to add value and better match the expected cashflows. In fact these supposed benefits are highly unlikely.

Some types of insurance-linked securities fit together more naturally. These include most cat ILS, whether they include the risk of natural catastrophes, manmade catastrophes or mortality. If the analytical expertise is available, other types can be added to this mix. But selectivity is key, as there are types of ILS that, from the point of view of the vast majority of investors, do not fit together well, often for reasons as simple as wide differences in duration and cashflow timing.

In addition, those who focus on delivering uncorrelated returns and minimising all risk in their portfolios with the exception of “pure” insurance risk do not want to pollute their portfolios with market risk present to a more significant degree in some other ILS classes.

There is, however, a small category of investment managers who can benefit from combining even very different types of insurance-linked securities in the same portfolio, and who have the expertise needed for the complex analysis and management on a portfolio basis. This could make sense only for the very few who can afford to deal with long and uncertain investment time horizons; and who possess the required combination of high-level analytical skills needed for managing such a portfolio.

- 1 We do not consider events so far in the tail of probability distribution that almost everything becomes correlated. For example, an earthquake of unprecedented magnitude leading to the loss of millions of lives in the developed world is such a tail event, but its probability is negligible.
- 2 See Chapter 5 for more information on the valuation of exchange-traded insurance derivatives.
- 3 Dividends, if they have been announced, are also reflected in valuation.

Conclusion

It is difficult to offer general conclusions after having examined an extensive list of insurance-linked securities that differ very significantly from each other. Despite the differences, however, they have sufficient characteristics in common to usefully be grouped together in the category of insurance-linked securities. The main common denominator is the insurance risk embedded in all of them.

PRACTITIONER'S VIEW

Discussion of the topics – individual securities, investment portfolio management, considerations of the investors and of (re)insurance companies or other hedgers, specific ILS structures, reinsurance, insurance risk modelling and all the other interrelated matters – has been undertaken from the point of view of a practitioner. The intent was to make the material useful in practical applications rather than to add to the body of existing academic research.

INSURANCE RISK

The book provides a detailed discussion of the products that bridge the gap between insurance/reinsurance and the capital markets. It discusses insurance risk – in its numerous forms – and how we can obtain investment exposure to this risk factor through insurance-linked securities, in order to take advantage of its exotic beta qualities. The investor point of view is the primary one, but the viewpoints of the insurance/reinsurance companies and the other parties involved in insurance securitisation are discussed as well. I hold the opinion that insurance and reinsurance companies, in their underwriting, also invest in insurance risk, albeit they do so in their own ways. This opinion is clearly expressed throughout the book, broadening the topic to discuss relevant reinsurance and insurance issues.

CURRENT STATE OF THE MARKET AND ITS FUTURE DEVELOPMENT

At this juncture, May 2010, it is not easy to easily predict in what direction the continuing convergence between the insurance/reinsurance and capital markets is going to take us. Throughout the book, moving from one product to another, I have attempted to make predictions of future developments. As sound as my logic may be, I know full well that some of the predictions I have made will likely be proved wrong. The one constant about this market is change. The experimentation never stops; new ideas and products appear and then sometimes disappear.

As a practitioner, I am both fascinated and incredibly frustrated by the uneven growth of the market. Innovation is the foundation of progress, but in business it should result in the development of new products and new markets that, once established, will grow along the developed path. The market exists and has expanded dramatically over the years. The fundamentals are in place to make this a much bigger market than it is right now, but the growth pattern will not be consistent. Not all of the products described in the book will see growth and some may even die out through natural selection. Others are here to stay and are firmly rooted for growth. The ILS market in general is certainly poised for expansion.

DRIVERS OF MARKET DEVELOPMENT

Looking back three main drivers of market development can be observed, that are applicable to most insurance-linked securities. These are: (1) insufficient capacity of the insurance and reinsurance industry to withstand the impact of catastrophic events without transferring some of the catastrophe risk to investors; (2) accounting rules that make it more capital-efficient to transfer some of the insurance risks to the capital markets; and (3) growing investor demand for assets that have low correlation with the traditional financial markets. These have been the key drivers and fundamental reasons for insurance securitisation. Continuing the list, other interrelated factors that can have a positive impact on the development of this market include (4) increasing emphasis on enterprise risk management that forces companies to better identify and manage their risks, including insurance risk; (5) growing realisation of the true magnitude of the unhedged insurance risk exposure, in particular to catastrophe risk; (6) focus on efficient capital management and shareholder value maximisation; (7) pressure from regulators and rating agencies to reduce catastrophe risk exposure; (8) continuous growth in the total amount of catastrophe insurance risk exposure; (9) improved modelling tools for quantification of insurance risk; (10)

development of expertise in insurance securitisation on the part of sponsors and issuers; (11) growth in expertise in insurance-linked securities in the investor community; (12) innovation in the form of developing new products and improved structures for existing ones; (13) movement towards greater transparency; and (14) movement towards standardisation.

The existence of virtually untapped sources of the insurance risk that is best borne by the capital markets is another reason for potential growth. Longevity is an example of such risk. In the short term, however, property catastrophe risk will remain the area of greatest activity.

OBSTACLES TO MARKET GROWTH

Some hold an overly optimistic view of the future of this market. In reality, there are many reasons why the market may go through difficult periods before its potential is realised.

The obstacles differ by product. General obstacles applicable to most insurance-linked securities include (1) insufficient understanding by potential hedgers of the magnitude of insurance risk they are holding; (2) unfamiliarity of the hedgers with the tools for transferring insurance risk to the capital markets; (3) inadequate investor interest in most types of insurance-linked securities; (4) imperfections of some securitisation structures; (5) high transaction costs; (6) basis risk concerns on the part of potential hedgers; (7) low average level of the understanding of insurance risk in the investor community; (8) inability of most investors to properly model insurance risk and other risks embedded in insurance-linked securities; (9) lack of confidence in the available insurance-risk modelling tools; (10) regulatory, accounting and tax concerns; (11) liquidity issues; (12) competition with other solutions to risk transfer such as reinsurance; (13) insufficient transparency; and (14) lack of standardisation. Simple inertia also belongs on this list.

Some solutions to overcoming these obstacles have in part already been implemented and are listed above as the factors contributing to the development of this market. Other solutions are product-specific as they have been designed to address the unique issues of these products.

OPPORTUNITIES

Insurance-linked securities have been a very inefficient market and most likely will stay this way for many years. While the inefficiency is an impediment to market growth, it also creates an opportunity for those who can explore it to their advantage through superior skill in assessing market

developments, sufficient flexibility to shift capital and resources to the most promising pockets of expected profitability, and a high level of expertise in the analysis of insurance risk. There is a window of opportunity for investors to enter the market and take advantage of the exotic beta potential offered by this asset class. Significant skill-based alpha may also be generated by those who possess superior expertise in insurance-linked securities. At the same time, the very reasons for this opportunity – market inefficiency and limited skill in the analysis of insurance-linked securities – present a danger and should suggest caution to those who do not have the required level of expertise.

Insurance and reinsurance companies have the opportunity to maximise shareholder value by incorporating insurance-linked securities in their arsenal of tools for managing risk and capital. In this process, like investors, the more sophisticated of them can use market inefficiencies to gain advantage over their competitors. However, the opportunities available to investors appear the most attractive. Not all of these opportunities are in insurance-linked securities that have originated directly with the insurance and reinsurance companies. Expertise in the analysis of insurance and related risks is key to taking advantage of these opportunities; this includes expertise both in the analysis of individual securities and in their management on a portfolio basis.

COMMENTS FROM READERS

I welcome any comments and suggestions from readers. I can be reached either through the publisher or at my personal e-mail address *alex.krutov@gmail.com*.

References

- Aase, K.**, 1999, "An Equilibrium Model of Catastrophe Insurance Futures and Spreads", *Geneva Papers on Risk and Insurance*, 24, pp. 69–96.
- Aase, K.**, 2001, "A Markov Model for the Pricing of Catastrophe Insurance Futures and Spreads", *The Journal of Risk and Insurance*, 68(1), pp. 25–50.
- ABI**, 2009, "The Financial Risks of Climate Change", ABI Research Paper No 19, Report Prepared by AIR Worldwide Corp. and the Met Office
- Achaerf, A.**, 2002, "Local Search Techniques for Constrained Portfolio Selection Problems", *Computational Economics*, 20(3), pp. 177–190.
- AIR Worldwide**, 2009, "The AIR Hurricane Model: AIR Atlantic Tropical Cyclone Model V11.0", Submitted to the Florida Commission on Hurricane Loss Projection Methodology, Final Submission (May).
- AIR Worldwide**, 2009, "AIR Worldwide Sponsors the Global Earthquake Model (GEM) Project", Press Release (February 25).
- Allen, F. and E. Carletti**, 2006, "Mark-to-Market Accounting and Liquidity Pricing", Working Paper 06–15, Wharton Financial Institutions Center.
- American Academy of Actuaries**, 1999, "Evaluating the Effectiveness of Index-Based Insurance Derivatives in Hedging Property/Casualty Insurance Transactions", Report of the Index Securitization Task Force, Washington, DC.
- American Academy of Actuaries**, 2002, Final Report of the American Academy of Actuaries Commissioners Standard Ordinary Task Force, Presented in Draft to the National Association of Insurance Commissioners Life and Health Actuarial Task Force, Washington, DC.
- Antolin, P. and H. Blommestein**, 2007, "Governments and the Market for Longevity-Indexed Bonds", Organisation for Economic Co-operation and Development Working Papers on Insurance and Private Pensions, OECD Publishing.
- Aon Benfield**, 2009, *Insurance-Linked Securities: Adopting to an Evolving Market*, (Chicago: Aon Benfield).
- Aon Benfield**, 2009, *Reinsurance Market Outlook*, (Chicago: Aon Benfield).
- Artzner P., F. Delbaen, J.-M. Eber and D. Heath**, 1999, "Coherent Measures of Risk", *Mathematical Finance*, 9(3), pp. 203–228.
- AXA Structured Finance**, 2009, "Facing the Storm with Catastrophe Bonds", *AXA Investment Manager's Research Review*, Paris: AXA.
- Barro, R. J.**, 2005, "What Price Catastrophe Risk?", *Business Week* (October 3).
- Bacon, C. R.**, 2008, *Practical Portfolio Performance Measurement and Attribution* (2nd Ed.), (New York: John Wiley & Sons).

- Baker, R.**, 1998, "Genetic Algorithms in Search and Optimization", *Financial Engineering News Magazine*, 2-3, pp. 1-13.
- Bakos, T. and K. Parankirinathan**, 2006, "The Life Settlement Market is an Opportunity", *The Journal of Structured Finance*, 12(2), pp. 46-49.
- Bantwal, V. J. and H. C. Kunreuther**, 2000, "A Cat Bond Premium Puzzle?", *Journal of Psychology and Financial Markets*, 1, pp. 76-91.
- Barbe, P., A.-L. Fougeres and C. Genest**, 2006, "On the Tail Behavior of Sums of Dependent Risks", *ASTIN Bulletin*, 36(2), pp. 361-373.
- Barnett, J., R. E. Kreps, J. A. Major and G. G. Venter**, 2007, "Multivariate Copulas for Financial Modeling", *Variance*, 1(1), pp. 103-119.
- Barrieu, P. and H. Louberge**, 2009, "Hybrid Cat Bonds," *The Journal of Risk and Insurance*, 76(3), pp. 547-578.
- Barrieu, P. and L. Albertini** (Eds), 2009, *The Handbook of Insurance-Linked Securities*, (New York: John Wiley & Sons).
- Beilis, A.**, 2000, "Weather Optionality", *Financial Engineering News Magazine*, 17.
- Bernstein, P. L.**, 1998, *Against the Gods: The Remarkable Story of Risk*, (New York: John Wiley & Sons).
- Bernstein, P. L.**, 2009, "Risk Inverse", *The Journal of Portfolio Management*, 35(3), pp. 1-1.
- AM Best Company**, 2005, "Life Settlement Securitization", www.ambest.com (September 1).
- AM Best Company**, 2006, "Assessing the 'Tail Risk' of Sidecars", www.ambest.com (October 9).
- AM Best Company**, 2007, "Securitization of Reinsurance Recoverables", www.ambest.com (August 20).
- AM Best Company**, 2008, "Rating Natural Catastrophe Bonds", www.ambest.com (January 22).
- AM Best Company**, 2008, "Securitization of Annuities", www.ambest.com (May 22).
- AM Best Company**, 2009, "Life Settlement Securitization", www.ambest.com (November 24).
- Bhuyan, V. B.**, Ed., 2009, *Life Markets*, (Hoboken, NJ: John Wiley & Sons).
- Biffis, E. and D. Blake**, 2008, "Securitizing and Tranching Longevity Exposures", Pensions Institute Discussion Paper PI-0824.
- Blake, D. and W. Burrows**, 2001, "Survivor Bonds: Helping to Hedge Mortality Risk", *The Journal of Risk and Insurance*, 68, pp. 339-348.
- Blake, D.**, 2006, *Pension Finance*, (New York: John Wiley & Sons).
- Blake, D., A. J. G. Cairns, K. Dowd and R. MacMinn**, 2006, "Longevity Bonds: Financial Engineering, Valuation and Hedging", *The Journal of Risk and Insurance*, 73, pp. 647-72.
- Blake, D., A. J. G. Cairns and K. Dowd**, 2006, "Living with Mortality: Longevity Bonds and Other Mortality-Linked Securities", *British Actuarial Journal*, 12, pp. 153-97.
- Blake, D., A. J. G. Cairns and K. Dowd**, 2008, "Longevity Risk and the Grim Reaper's Toxic Tail: The Survivor Fan Charts", *Insurance: Mathematics & Economics*, 42, pp. 1062-66.
- Blake, D. and D. Harrison**, 2008, "Life Settlements Trade Raises Complex Issues", *Financial Times* (June 6).
- Blake, D., T. Boardman, A. Cairns and K. Dowd**, 2009, "Everyone Wins if UK Issues Longevity Bond", *Financial Times* (June 28).
- Bodoff, N. M. and Y. Gan**, 2009, "An Analysis of the Market Price of Cat Bonds", *Casualty Actuarial Society Forum*, Spring, pp. 1-26.

- Booth, H., R. J. Hyndman, L. Tickle and P. de Jong, 2006, "Lee-Carter Mortality Forecasting: A Multi-Country Comparison of Variants and Extensions", *Demographic Research*, 15, pp. 289–310.
- Booth, H. and L. Tickle, 2008, "Mortality Modelling and Forecasting: A Review of Methods", *Annals of Actuarial Science*, 3, pp. 3–43.
- Borak, S., W. Härdle and R. Weron, 2005, *Statistical Tools for Finance and Insurance*, (Berlin: Springer).
- Bouriaux, S., 2001, "Basis Risk, Credit Risk and Collateralization Issues for Insurance-Linked Derivatives and Securities", *Journal of Insurance Regulation*, 20(1), pp. 94–120.
- Bouriaux, S. and R. MacMinn, 2009, "Securitization of Catastrophe Risk: New Developments in Insurance-Linked Securities and Derivatives", *Journal of Insurance Issues*, 32(1), pp. 1–34.
- Bowers, N., H. Gerber, J. Hickman, D. Jones and C. Nesbitt, 1997, *Actuarial Mathematics*, (Schaumburg, IL: Society of Actuaries).
- Boyle, P., M. Hardy and T. C. F. Vorst, 2005, "Life after VaR", *The Journal of Derivatives*, 13(1), pp. 48–55.
- Boyle, P. and W. Tian, 2007, "Portfolio Management with Constraints", *Mathematical Finance*, 17(3), pp. 319–344.
- Brabazon, A., and M. O'Neill, 2006, *Biologically Inspired Algorithms for Financial Modeling*, (Berlin: Springer-Verlag).
- Brady, M., 2008, "How Will Longer Life Expectancy Estimates Impact Settlements?", *National Underwriter* (December 3).
- Brockett, P., M. Wang, C. Yang and H. Zou, 2006, "Portfolio Effects and Valuation of Weather Derivatives", *Financial Review*, 41, pp. 55–76.
- Brockett, P., L. Golden, M.-M. Wen and C. C. Yang, 2009, "Pricing Weather Derivatives Using the Indifference Pricing Approach", *North American Actuarial Journal*, 13(3), pp. 303–315.
- Bromann, K., 2008, "ILS Investments and Portfolio Diversification", Presentation, Workshop on Insurance-Linked Securities, Imperial College, London (October 31).
- Cain, M. and D. Peel, 2004, "Utility and the Skewness of Return in Gambling", *Geneva Papers on Risk and Insurance*, 29(2), pp. 145–163.
- Cairns, A. J. G., D. Blake and K. Dowd, 2004, "Pricing Framework for Securitization of Mortality Risk", Technical Report, Heriot-Watt University.
- Cairns, A. J. G., D. Blake and K. Dowd, 2006, "Pricing Death: Frameworks for the Valuation and Securitization of Mortality Risk", *ASTIN Bulletin*, 36, pp. 79–120.
- Cairns, A., D. Blake, D. and K. Dowd, 2006, "A Two-Factor Model for Stochastic Mortality with Parameter Uncertainty: Theory and Calibration", *The Journal of Risk and Insurance*, 73, pp. 687–718.
- Cairns, A. J. G., D. Blake and K. Dowd, 2008, "Modelling and Management of Mortality Risk: A Review", Pension Institute Discussion Paper PI-0814.
- Cairns, A. J. G., D. Blake, K. Dowd, G. D. Coughlan, and M. Khalaf-Allah, 2008, "Mortality Density Forecasts: An Analysis of Six Stochastic Mortality Models", Pension Institute Discussion Paper PI-0801.
- Cairns, A. J. G., D. Blake, K. Dowd, G. D. Coughlan, D. Epstein, A. Ong, and I. Balevich, 2009, "A Quantitative Comparison of Stochastic Mortality Models Using Data from England and Wales and the United States", *North American Actuarial Journal*, 13(1), pp. 1–35.
- Cairns, A., 2009, "The Government is Planning to Raise the Age of Retirement ... but are Scots Getting a Fair Deal?", *The Sunday Herald* (September 20).

- Campbell, K. and W. Keogh, 2009, "Understanding of Earthquake Risk Improving Dramatically in U.S.", EQECAT, ABS Consulting Group.
- Canabarro E., M. Finkemeier, R. R. Anderson and F. Bendimerad, 2000, "Analyzing Insurance-Linked Securities", *The Journal of Risk Finance*, 1(2), pp. 49–75.
- Canter, M. S., J. B. Cole and R. L. Sandor, 1996, "Insurance Derivatives: A New Asset Class for the Capital Markets and a New Hedging Tool for the Insurance Industry", *Journal of Derivatives*, 4, pp. 89–105.
- Cao, M. and J. Wei, 2004, "Weather Derivatives Valuation and Market Price of Weather Risk", *Journal of Futures Markets*, 24(11), pp. 1065–1089.
- Casey, B. T. and T. D. Sherman, 2007, "Are Life Settlements a Security?", *The Journal of Structured Finance*, 12(4), pp. 55–60.
- CFA Institute, 2006, *Global Investment Performance Standards (GIPS) Handbook* (2nd Ed.), (Charlottesville, VA: CFA Institute).
- Chacko, G., P. Hecht, V. Dessain, A. Sjöman and A. J. Plotkin, 2004, "Bank Leu's Prima Cat Bond Fund", Harvard Business School.
- Chance, D., 2004, "Default Risk as an Option", *Financial Engineering News Magazine*, 38 (January/February), pp. 15–22.
- Chen, H. and S. Cox, 2009, "Modeling Mortality with Jumps: Applications to Mortality Securitization", *The Journal of Risk and Insurance*, 76(3), pp. 727–751.
- Chen, H. and J. D. Cummins, 2009, "Longevity Bond Premiums: The Extreme Value Approach and Risk Cubic Pricing", *Insurance: Mathematics and Economics*, Submitted.
- Chen, S.-H., 1998, "Evolutionary Computation in Financial Engineering: A Road Map of GAs and GP", *Financial Engineering News Magazine*, 6, pp. 3–11.
- Chua, D. B., M. Kritzman and S. Page, 2009, "The Myth of Diversification", *The Journal of Portfolio Management*, 36(1), pp. 26–35.
- Chung, J., 2007, "Single Product Sector Urged for Longevity Risk", *Financial Times* (July 2).
- Clarke, R. G., H. de Silva and B. Wander, 2002, "Risk Allocation versus Asset Allocation", *The Journal of Portfolio Management*, 29, pp. 9–30.
- Cobley, M., 2008, "U.K. Firms Face Rising Longevity of Workers", *The Wall Street Journal* (July 1).
- Cohen, N. and J. Lemer, 2009, "Babcock Pension to Hedge Risk of Longevity", *Financial Times* (May 13).
- Cole, C. R. and K. A. McCullough, 2006, "A Reexamination of the Corporate Demand for Reinsurance", *The Journal of Risk and Insurance*, 73(1), pp. 169–192.
- Congressional Budget Office, 2005, "A Potential Influenza Pandemic: Possible Macroeconomic Effects and Policy Issues", CBO (December 8).
- Conning Research and Consulting, Inc., 2007, *Life Settlement Market: Increasing Capital and Investor Demand*, (Hartford, CT: Conning Research and Consulting).
- Conning Research and Consulting, Inc., 2008, *Life Settlements: New Challenges to Growth*, (Hartford, CT: Conning Research and Consulting).
- Cossette, H., T. Duchesne and E. Marceau, 2003, "Modeling Catastrophes and Their Impact on Insurance Portfolios", *North American Actuarial Journal*, 7(4), pp. 1–22.
- Coughlan, G., D. Epstein, A. Ong, A. Sinha, J. Hevia-Portocarrero, E. Gingrich, M. Khalaf-Allah and P. Joseph, 2007, *LifeMetrics: A Toolkit for Measuring and Managing Longevity and Mortality Risks*, Technical Document, JP Morgan.

- Coughlan, G., D. Epstein, A. Sinha and P. Honig, 2007, "q-Forwards: Derivatives for Transferring Longevity and Mortality Risk", JP Morgan Pension Advisory Group, JP Morgan.
- Cowley, A. and J. D. Cummins, 2005, "Securitization of Life Insurance Assets and Liabilities", *The Journal of Risk and Insurance*, 72, pp. 193–226.
- Cox, S. H. and H. W. Pedersen, 2000, "Catastrophe Risk Bonds", *North American Actuarial Journal*, 4(4), pp. 56–82.
- Cox, S. H. and Y. Lin, 2004, "Natural Hedging of Life and Annuity Mortality Risks", *Proceedings of the 14th International AFIR Colloquium*, pp. 483–507.
- Cox, S. H., Y. Lin and S. Wang, 2006, "Multivariate Exponential Tilting And Pricing Implications For Mortality Securitization", *The Journal of Risk and Insurance*, December, 73(4), pp. 719–736.
- Cox, S. and Y. Lin, 2007, "Natural Hedging of Life and Annuity Mortality Risks", *North American Actuarial Journal*, 11(3), pp. 1–15.
- Crama, Y. and M. Schyns, 2003, "Simulated Annealing for Complex Portfolio Selection Problems", *European Journal of Operational Research*, 150(3), pp. 546–571.
- Cremers, J. H., M. Kritzman and S. Page, 2005, "Optimal Hedge Fund Allocations", *The Journal of Portfolio Management*, 31(3), pp. 70–81.
- Csiszar, E. N., 2007, "An Update on the Use of Modern Financial Instruments in the Insurance Sector", *The Geneva Papers on Risk and Insurance*, 32, pp. 319–331.
- Culp, C. L., 2006, *Structured Finance and Insurance*, (Hoboken, NJ: Wiley & Sons).
- Cummins, J. D., D. Lalonde and R. D. Phillips, 2004, "The Basis Risk of Index-Linked Catastrophic Loss Securities", *Journal of Financial Economics*, 71, pp. 77–111.
- Cummins, J. D., 2005, "Convergence in Wholesale Financial Services: Reinsurance and Investment Banking", *The Geneva Papers on Risk and Insurance*, 30, pp. 187–22.
- Cummins, J. D., 2006, "Should the Government Provide Insurance for Catastrophes?", *Federal Reserve Bank of St. Louis Review*, 88, pp. 337–379.
- Cummins, J. D., 2008, "Cat Bond and Other Risk-Linked Securities: State of the Market and Recent Developments", *Risk Management and Insurance Review*, 11(1), pp. 23–47.
- Cummins, J. D. and P. Trainar, 2009, "Securitization, Insurance and Reinsurance", *The Journal of Risk and Insurance*, 76(3), pp. 463–492.
- Cummins, J. D. and M. A. Weiss, 2009, "Convergence of Insurance and Financial Markets: Hybrid and Securitized Risk-Transfer Solutions", *The Journal of Risk and Insurance*, 76(3), pp. 493–545.
- Currie, I. D., M. Durban and P. H. C. Eilers, 2004, "Smoothing and Forecasting Mortality Rates", *Statistical Modelling*, 4, pp. 279–98.
- Dahl, M., M. Melchior and T. Muller, 2008, "On Systematic Mortality Risk and Risk Minimisation with Survivor Swaps", *Scandinavian Actuarial Journal*, 2–3, pp. 114–46.
- Davies, J., 2007, "Swiss Re Indices Seek to Make Catastrophe Clearer", *Financial Times* (July 2)
- Davies, J., 2009, "Opportunities in Risk", *Financial Times* (October 16).
- Dawson, P., D. Blake, A. J. G. Cairns and K. Dowd, 2007, "Completing the Survivor Derivatives Market", Pensions Institute Discussion Paper PI-0712.
- DBRS, 2008, "Increased Life Expectance: Implications for Life Settlement Transactions", *US Structured Finance Newsletter*, 4(25) (June 23).
- DBRS, 2008, "Rating U.S. Life Settlement Securitizations", New York (February).

- Deutsche Börse**, 2008, "Deutsche Börse Launches Business with Longevity Data", Press Release (March 11).
- Deutsche Börse**, 2008, "Deutsche Börse Launches Longevity Indices", Press Release (December 1).
- Dhillon, H.**, 2009, "Winds of Change for Cat Bond Market", *Credit* (June).
- Doherty, N.**, 1997, "Financial Innovations in the Management of Catastrophic Risk", *Journal of Applied Corporate Finance*, 10, pp. 84–95.
- Doherty, N. A. and H. J. Singer**, 2002, "The Benefits of a Secondary Market for Life Insurance Policies", The Wharton Financial Institutions Center.
- Dorr, D. C.**, 2007, "Longevity Trading: Bridging the Gap between the Insurance Markets and the Capital Markets", *Journal of Structured Finance*, 13(2), pp. 50–53.
- Dowd, K.**, 2005, "Coherent Risk Measures", *Financial Engineering News Magazine*, 41, pp. 9–10.
- Dowd, K.**, 2005, "Distortion Risk Measures", *Financial Engineering News Magazine*, 44, pp. 7–14.
- Dowd, K., A. J. G. Cairns and D. Blake**, 2006, "Mortality-Dependent Financial Risk Measures", *Insurance: Mathematics and Economics*, 38, pp. 427–40.
- Dowd, K., D. Blake, A. J. G. Cairns and P. Dawson**, 2006, "Survivor Swaps", *The Journal of Risk and Insurance*, 73(1), pp. 1–17.
- Dowd, K.**, 2006, "The Invisible Problem of Risk Blindness", *Financial Engineering News Magazine*, 52, pp. 5–10.
- Dowd, K., D. Blake and A. J. G. Cairns**, 2006, "The Grave Problem of Longevity Risk", *Financial Engineering News Magazine*, 49, pp. 19, 30.
- Dowd, K., A. J. G. Cairns, D. Blake, G. D. Coughlan, D. Epstein and M. Khalaf-Allah**, 2008, "Evaluating the Goodness of Fit of Stochastic Mortality Models", Pensions Institute Discussion Paper PI-0802.
- Durbin, D.**, 2001, "Managing Natural Catastrophe Risks: The Structure and Dynamics of Reinsurance", *The Geneva Papers on Risk and Insurance*, 26, pp. 297–309.
- Eling, M. and D. Toplek**, 2009, "Modeling and Management of Nonlinear Dependencies – Copulas in Dynamic Financial Analysis", *The Journal of Risk and Insurance*, 76(3), pp. 651–681.
- Embrechts, P., C. Klüppelberg and T. Mikosch**, 2004, *Modelling Extremal Events: For Insurance and Finance*, (Berlin: Springer).
- Embrechts, P., J. Neslehova and M. V. Wüthrich**, 2009, "Additivity Properties for Value-at-Risk under Archimedean Dependence and Heavy-Tailedness", *Insurance: Mathematics and Economics*, 44(2), pp. 164–169.
- EQECAT**, 2008, "eCAT: A New Tool for Managing Securitized Natural Catastrophe Risk", ABS Group.
- EQECAT**, 2009, "Using Recent Hurricane Data to Evaluate Prospective Risk Models", ABS Group.
- EQECAT**, 2009, EQECAT Florida Hurricane Model 2009, Submitted to the Florida Commission on Hurricane Loss Projection Methodology (Revised May 18).
- Eskandari, H. and C. D. Geiger**, 2008, "A Fast Pareto Genetic Algorithm Approach for Solving Expensive Multiobjective Optimization Problems", *Journal of Heuristics*, 14(3), pp. 203–241.
- Estrada, J.**, 2006, "Downside Risk in Practice", *Journal of Applied Corporate Finance*, 18(1), pp. 117–125.
- Evans, C.**, 2009, "Norwich Union Completes 475 Million Pounds Longevity Swap", Thompson Reuters (March 24).

- Evans, C., 2009, "Catastrophe Derivatives Demand Surges as Storm Season Nears", *Insurance Journal*, (May 19).
- Felsted, A., 2008, "RMS Launches European Wind Storm Risk Index", *Financial Times* (February 15).
- Finkelstein, M., H. G. Tucker, and J. A. Veeh, 2006, "Pareto Tail Index Estimation Revisited", *North American Actuarial Journal*, 10(1), pp. 1–10.
- Fitch, 2008, "Insurance-Linked Securities Ratings Criteria", Fitch Ratings (February 4).
- Fitch, 2008, "Catastrophe Bonds: Ratings Criteria", Fitch Ratings (March 11).
- Fitch, 2008, "Life Insurance Reserve Finance: Ratings Criteria", Fitch Ratings (July 8).
- Fitch, 2008, "Zest: AEGON's Scottish Equitable Value of In-Force Securitisation", Fitch Ratings (July 22).
- Freeman, M. B., 2007, "Lurking Pitfalls", *The Journal of Structured Finance*, 13(2), pp. 83–86.
- Froot, K. A. and M. S. Seasholes, 1997, "USAA: Catastrophe Risk Financing", Harvard Business School.
- Froot, K. A., Ed., 1999, *The Financing of Catastrophe Risk*, (The University of Chicago Press).
- Froot, K. A., 2001, "The Market for Catastrophe Risk: A Clinical Examination", *Journal of Financial Economics*, 60, pp. 529–571.
- Froot, K. A., 2007, "Risk Management, Capital Budgeting, and Capital Structure Policy of Insurers and Reinsurers", *The Journal of Risk and Insurance*, 74, pp. 273–299.
- Froot, K. A. and P. O'Connel, 2008, "On the Pricing of Intermediated Risks: Theory and Application to Catastrophe Reinsurance", *Journal of Banking and Finance*, 32, pp. 69–85.
- Fulla, S. and J.-P. Laurent, 2008, "Mortality Fluctuations Modelling with a Shared Frailty Approach", *Life and Pension Risk* (October).
- Gehin, W. M. and M. Vaissie, 2006, "The Right Place for Alternative Betas in Hedge Fund Performance: An Answer to the Capacity Effect Fantasy", *The Journal of Alternative Investments*, 9, pp. 9–18.
- Geman, H., Ed., 1999, *Insurance and Weather Derivatives: From Exotic Options to Exotic Underlyings*, (London: Risk Books).
- General Accounting Office, 2002, "Catastrophe Insurance Risks: The Role of Risk-Linked Securities and Factors Affecting Their Use" (GAO-02-941), Report to the Chairman, Committee on Financial Services, US House of Representatives, Washington, DC.
- Gilli, M., E. Kellezi and H. Hysi, 2006, "A Data-Driven Optimization Heuristic for Downside Risk Minimization", *The Journal of Risk*, 8(3), pp. 1–18.
- Glassermann, P., 2004, *Monte Carlo Methods in Financial Engineering*, (New York: Springer).
- Golden, L., M. Wang and C. Yang, 2007, "Handling Weather Related Risks Through Financial Markets: Considerations of Credit Risk, Basis Risk, and Hedging", *The Journal of Risk and Insurance*, 74(2), pp. 319–46.
- Goldstein, M., 2007, "Profiting from Mortality", *Business Week* (October 30).
- Goldstein, M., 2008, "Why Death Bonds Look So Frail", *Business Week* (February 25).
- Gore, G., 2006, "Credit Suisse Launched Longevity Index", *Risk*, 19(1).
- Gorvett, R. W., 1999, "Insurance Securitization: The Development of a New Asset Class", Securitization of Risk – Discussion Paper Program, Casualty Actuarial Society, pp. 133–173.
- Government Accountability Office, 2005, "Catastrophe Risk: U.S. and European Approaches to Insure Natural Catastrophe and Terrorism Risks" (GAO-05-1999), Report to the Chairman, Committee on Financial Services, US House of Representatives, Washington, DC.

- Group of Thirty**, 2006, *Reinsurance and International Financial Markets*, Washington, DC.
- Grossi, P. and H. Kunreuther**, 2004, *Catastrophe Modeling: A New Approach to Managing Risk*, (New York: Springer).
- Gurenko, E. N. (Ed.)**, 2004, *Catastrophe Risk and Reinsurance*, (London: Risk Books).
- Guy Carpenter**, 2008, "2008 Reinsurance Market Review", New York.
- Guy Carpenter**, 2008, "The Catastrophe Bond Market at Year-End 2007", New York.
- Hardy, M. R. and J. L. Wirth**, 2004, "The Iterated CTE – A Dynamic Risk Measure", *North American Actuarial Journal*, 8(4), pp. 62–75.
- Hardy, M. R.**, 2005, "A Matter of Life and Death", *Financial Engineering News Magazine*, 46, pp. 17–20.
- Harrison, J.**, 2007, "Securitisation of Insurance Risks", *Financial Instruments Tax and Accounting Review*, September, pp. 4–8.
- Hill, A.**, 2009, "Longevity Risk Pioneers Must Learn from Crisis", *Financial Times* (May 12).
- Hill, J. M.**, 2006, "Alpha as a Net Zero-Sum Game", *Journal of Portfolio Management*, 32, pp. 24–32.
- Horsewood, R.**, 2003, "A Farewell to Arms: What the DARPA Funded Market in Geopolitical Risk Futures Would Have Looked Like", *Financial Engineering News Magazine*, 34 (November/December).
- Hull, J., M. Predescu and A. White**, 2005, "Bond Prices, Default Probabilities and Risk Premiums", *The Journal of Credit Risk*, 1(2), pp. 53–60.
- Hull, J. C.**, 2006, *Options, Futures, and Other Derivatives* (6th Ed.), (Upper Saddle River, NJ: Prentice Hall).
- IAIS**, 2009, "Developments in (Re)Insurance Securitisation, Global Reinsurance Market Report", International Association of Insurance Supervisors, Basel (August 26).
- International Financing Review**, 2006. "Mortality Bond Issue Upsized", 40 (November 4).
- International Monetary Fund**, 2006, *The Limits of Market-based Risk Transfer and Implications for Managing Systemic Risks*, Washington, DC.
- Johnson, K.**, 2008, "Stormy Weather: Is Global Warming to Blame?", *The Wall Street Journal* (August 4).
- Johnson, S.**, 2007, "FSA Warns over Longevity Figures", *Financial Times* (May 14).
- Jones, S.**, 2009, "The Formula that Felled Wall Street", *Financial Times* (April 24).
- Jorion, P.**, 2007, *Value at Risk: The New Benchmark for Managing Financial Risk* (3rd Ed.), (New York: McGraw-Hill).
- Jung, J.**, 2008, "Living with Volatility", *Risk Magazine* (March)
- Juri, A. and M. V. Wutrich**, 2003, "Tail Dependence from a Distributional Point of View", *Extremes*, 6(3), pp. 213–246.
- Kahneman, D. and D. Lovello**, 1993, "Timid Choices and Bold Forecasts: A Cognitive Perspective on Risk Taking", *Management Science*, 39, pp. 17–31.
- Keenlyside, N., M. Latif, J. Junclaus, L. Kornblueh and E. Roeckner**, 2008, "Advancing Decadal Climate Scale Prediction in the North Atlantic", *Nature*, 453, pp. 84–88.
- Klein, R. W. and S. Wang**, 2009, "Catastrophe Risk Financing in the United States and the European Union: A Comparative Analysis of Alternative Regulatory Approaches", *The Journal of Risk and Insurance*, 76(3), pp. 607–637.

- Klotzbach, P. J. and W. M. Gray, 2009, "Twenty-Five Years of Atlantic Basin Seasonal Hurricane Forecasts", *Geophysical Research Letters*, 36(9), L09711.
- Kothari, V., 2006, *Securitization: The Financial Instrument of the Future*, (Singapore: John Wiley & Sons).
- Krokhmal, P., J. Palmquist and S. Uryasev, 2002, "Portfolio Optimization with Conditional Value-at-Risk Objective and Constraints", *The Journal of Risk*, 4(2), pp. 11–27.
- Krutoy, A., 2006, "Insurance-Linked Securities: An Emerging Class of Financial Instruments", *Financial Engineering News Magazine*, 48, pp. 7–16.
- Krutoy, A., 2006, "Insurance Derivatives: A Hidden Market", *Financial Engineering News Magazine*, 51, pp. 8–12.
- Kunreuther, H., 2002, "The Role of Insurance in Managing Extreme Events: Implications for Terrorism Coverage", *Business Economics* (April).
- Lane, M., "Pricing Risk Transfer Transactions", *ASTIN Bulletin* 30:2, 2000, pp. 259–293.
- Lane, M. (Ed.), 2002, *Alternative Risk Strategies*, (London: Risk Books).
- Lane, M. and R. Beckwith, 2007, "Developing LFC Return Indices for Insurance Securitizations", Lane Financial.
- Lane, M., 2007, "Of Sidecars and Such", Lane Financial.
- Lane, M. and O. Mahul, 2008, "Catastrophe Risk Pricing: An Empirical Analysis", *Policy Research Working Paper Series*, The World Bank (November 1).
- Lane, M., 2009, "Genuine Alpha, Perfect Security – Reaffirming ILS Rationales", Lane Financial.
- LaRow, T. E., Y.-K. Lim, D. W. Shin, E. P. Chassignet and S. Cocke, 2008, "Atlantic Basin Seasonal Hurricane Simulations", *Journal of Climate*, 21, pp. 3191–3206.
- Laster, D. and M. Raturi, 2001, "Capital Market Innovation in the Insurance Industry", *sigma* no. 3/2001, Swiss Re, Zurich.
- Lee, B., 2000, "A Consistent Approach to Measuring Hedge Effectiveness", *Financial Engineering News Magazine*, 14 (February), pp. 5–13.
- Lee, J. P. and M. T. Yu, 2002, "Pricing Default-Risky Cat Bonds with Moral Hazard and Basis Risk", *The Journal of Risk and Insurance*, 69, pp. 25–44.
- Lee, R. D. and L. Carter, 1992, "Modelling and Forecasting the Time Series of U.S. Mortality", *Journal of the American Statistical Association*, 87, pp. 659–71.
- Lee, R. D., 2000, "The Lee-Carter Method of Forecasting Mortality, with Various Extensions and Applications", *North American Actuarial Journal*, 4, pp. 80–93.
- Leibowitz, M. L., 2005, "Alpha Hunters and Beta Grazers," *Financial Analysts Journal*, 61, pp. 32–39.
- Lewis, M., 2007, "In Nature's Casino", *New York Times Magazine* (August 26).
- Li, J. and S. Taiwo, 2006, "Enhancing Financial Decision Making Using Multi-Objective Financial Genetic Programming", *Proceedings of IEEE Congress on Evolutionary Computation*, pp. 2171–2178 (July 16–21).
- Li, N., R. Lee and S. Tuljapurkar, 2004, "Using the Lee-Carter Method to Forecast Mortality for Populations with Limited Data", *International Statistical Review*, 72, pp. 19–36.
- Lin, D. and S. Wang, 2002, "A Genetic Algorithm for Portfolio Selection Problems", *Advanced Modeling and Optimization*, 4(1), pp. 13–27.
- Lin, X. S. and X. Liu, 2007, "Markov Ageing Process and Phase-Type Law of Mortality", *North American Actuarial Journal*, 11(4), pp. 92–109.

- Lin, Y. and S. H. Cox**, 2005, "Securitization of Mortality Risks in Life Annuities", *The Journal of Risk and Insurance*, 72, pp. 227–252.
- Liu, J.**, 2007, "Portfolio Selection in Stochastic Environment", *The Review of Financial Studies*, 20(1), pp.1–39.
- Litterman, R.**, 2005, "Active Alpha Investing", Goldman Sachs Asset Management, New York.
- Loeys, J., N. Panigirtzoglou and R. Ribeiro**, 2007, "Longevity: A Market in the Making", Global Market Strategy, JP Morgan.
- Los, C. A.**, 2005, "When to Put All Your Eggs in One Basket? When Diversification Increases Portfolio Risk!", *Financial Engineering News Magazine*, 41 (January-February), pp. 4–9.
- Louberge, H., E. Kellezi and M. Gilli**, 1999, "Using Catastrophe-Linked Securities to Diversify Insurance Risk: A Financial Analysis of Cat Bonds", *Journal of Insurance Issues*, 22, pp. 125–146.
- Lyon, P.**, 2003, "Securitising Terror", *Risk Magazine* (March).
- Lucida**, 2008, "Lucida and JPMorgan First to Trade Longevity Derivative", Press Release (15 February).
- Maginn, J. L. et al** (Eds.), 2007, *Managing Investment Portfolios: A Dynamic Process* (3rd Ed.), CFA Institute, (Hoboken, NJ: John Wiley & Sons).
- Markowitz, H.**, 1952, "Portfolio Selection", *Journal of Finance*, 7(1), pp. 77–91.
- Matthys, G., E. Delafosse, A. Guillou and J. Beirlant**, 2004, "Estimating Catastrophic Quantile Levels for Heavy-Tailed Distributions", *Insurance: Mathematics and Economics*, 34, pp. 517–537.
- McNeil, A., R. Frey, and P. Embrechts**, 2005, *Quantitative Risk Management: Concepts, Techniques, and Tools*. (Princeton University Press).
- McRaith, M. T.**, 2009, "Life Settlements and the Need for Increasing Transparency", Testimony Before the United States Senate Special Committee on Aging (April 29).
- Meyers, G. G.**, 2007, "The Common Shock Model for Correlated Insurance Losses", *Variance*, 1(1), pp. 40–52.
- Mijuk, G.**, 2008, "Port in the Storm: Insurers' Catastrophe Bonds", *The Wall Street Journal* (October 14).
- Milevsky, M. A. and S. D. Promislow**, 2001, "Mortality Derivatives and the Option to Annuitize", *Insurance Mathematics and Economics*, 29(3), pp. 299–318.
- Modu E.**, 2004, "Life Settlement Securitization", *A.M. Best Report* (October 18).
- Modu, E.**, 2006, "Gauging the Basis Risk of Catastrophe Bonds", *A.M. Best Report* (September 25).
- Moody's**, 2007, "Reinsurance Sidecars: Moody's Five Principles", *Moody's* (March 28).
- JP Morgan Chase and Co.**, 2007, "JP Morgan Launches Longevity Index: Investment bank Creates Life-Metrics Platform", Press Release (March 13).
- Mott, A. R.**, 2007, "New Swaps to Hedge Alpha and Beta Longevity Risks of Life Settlement Pools", *Journal of Structured Finance*, Summer, pp. 54–61.
- Muermann, A.**, 2001, "Pricing Catastrophe Insurance Derivatives", Working Paper, London School of Economics.
- Muermann, A.**, 2003, "Actuarially Consistent Valuation of Catastrophe Derivatives", *The Wharton Financial Institution Center Working Paper Series*, 03–18, University of Pennsylvania.
- Mutenga, S. and S. Staikouras**, 2007, "The Theory of Catastrophe Risk Financing: A Look at the Instruments that Might Transform the Insurance Industry", *The Geneva Papers on Risk and Insurance*, 32, pp. 222–45.

- Muir-Wood, R.**, 2006, "RMS Advances the Science of Catastrophe Risk Models", *Bermuda Re*, pp. 1–4 (February).
- Nakada, P.**, 2009, "Insurance-Linked Securities: Last Asset Class Standing", *Insurance, Finance & Investment*, XIV(6).
- Natarajan, K., D. Pachamanova and M. Sim**, 2008, "Incorporating Asymmetric Distributional Information in Robust Value-at-Risk Optimization", *Management Science*, 54(3), pp. 573–585.
- Nell, M. and A. Richter**, 2004, "Improving Risk Allocation through Indexed Cat Bonds", *The Geneva Papers on Risk and Insurance*, 29 (2), pp. 183–201.
- Niehaus, G.**, 2002, "The Allocation of Catastrophe Risk", *Journal of Banking and Finance*, 26, pp. 585–596.
- Nott, J., J. Haig, H. Neil, and D. Gillieson**, 2007, "Greater Frequency Variability of Landfalling Tropical Cyclones at Centennial Compared to Seasonal and Decadal Scales", *Earth and Planetary Science Letters*, 255(3–4), pp. 367–372.
- Palmer, T. N., F. J. Doblas-Reyes, A. Weisheimer and M. J. Rodwell**, 2008, "Toward Seamless Prediction Calibration of Climate Change Projections Using Seasonal Forecasts", *Bulletin of the American Meteorological Society*, 89, pp. 459–470.
- Papachristou, D.**, 2009, "Statistical Analysis of Spreads at the Time of Issue", ASTIN Colloquium, Helsinki (June).
- Patel, N.**, 2007, "For Whom Cat Risk Tolls", *Risk Magazine* (November).
- Pelsser, A.**, 2008, "On the Applicability of the Wang Transform for Pricing Financial Risks", *ASTIN Bulletin*, 38(1), pp. 171–181.
- Pengelly, M.**, 2008, "Second Life", *Risk Magazine* (March).
- Perera, N. and L. Pearson**, 2007, "An Exploration of Mortality Risk Mitigation", *The Journal of Structured Finance*, 13(2), pp. 44–49.
- Perez-Fructuoso, M. J.**, 2008, "Modeling Loss Index Triggers for Cat Bonds: A Continuous Approach", *Variance* 2(2), pp. 253–265.
- Plantin, G.**, 2006, "Does Reinsurance Need Reinsurers?", *The Journal of Risk and Insurance*, 73(1), pp. 153–168.
- Poncet, P. and V. Vaugirard**, 2002, "The Pricing of Insurance-Linked Securities under Interest Rate Uncertainty", *The Journal of Risk Finance*, 3(3), pp. 48–59.
- Potter, N. F. and M. McDonnell**, 2008, "Life Insurance Securitizations: Legal Structures for Statutory Reserve Financings and Embedded Value Monetizations", *The Journal of Structured Finance*, 4(3), pp. 75–79.
- Price, K., Storn, R. and J. Lampinen**, 2005, *Differential Evolution: A Practical Approach to Global Optimization*, (Berlin: Springer).
- Pugh, H.**, 2005, "Pandemic: The Cost of Avian Influenza", *Contingencies*, September–October, pp. 22–27.
- Reactions**, 2004, "Make Money Rain or Shine" (Case Study: Coriolis Capital), *Reactions*, November, pp. 48–50.
- Reinstein, A. and C. Miller**, 2007, "Accounting for the Purchase of Life Settlement Contracts", *The CPA Journal*, 77(9), pp. 28–34.
- Renshaw, A. E. and S. Haberman**, 2003, "Lee-Carter Mortality Forecasting with Age-Specific Enhancement", *Insurance: Mathematics and Economics*, 33, pp. 255–272.
- Renshaw, A. E. and S. Haberman**, 2006, "A Cohort-Based Extension to the Lee-Carter Model for Mortality Reduction Factors", *Insurance: Mathematics and Economics*, 38, pp. 556–570.

- Richard, C.**, 2004, "With \$70M Bond Deal, Wall Street Manages to Securitize Death", *The Wall Street Journal* (April 30).
- Rind, K. and P. Siegert**, 2008, *Taxation of Life Insurance Policies in an Evolving Secondary Marketplace*, (Keystone, CO: Insurance Studies Institute).
- RMS**, 2005, "Hurricane Katrina: Profile of a Super Cat", Risk Management Solutions, Inc.
- RMS**, 2007, "RMS Pandemic Influenza Model Overview", Risk Management Solutions, Inc.
- RMS**, 2009, "Miu Platform", Risk Management Solutions, Inc.
- RMS**, 2009, "Paradex", Risk Management Solutions, Inc.
- RMS**, 2009, "RMS US Hurricane Model", Submitted to the Florida Commission on Hurricane Loss Projection Methodology (Revised May 22).
- Rockafellar, R. T. and S. P. Uryasev**, 2000, "Optimization of Conditional Value-at-Risk", *The Journal of Risk*, 2(3), pp. 21–42.
- Rockafellar, R. T. and S. P. Uryasev**, 2002, "Conditional Value-at-Risk for General Loss Distribution", *Journal of Banking and Finance*, 26, pp. 1443–1471.
- Ross, S.**, 2004, "Compensation, Incentives, and the Duality of Risk Aversion and Riskiness", *Journal of Finance*, 59, pp. 207–25.
- Rundle, J. B., P. B. Rundle, A. Donnellan, D. L. Turcotte et al**, 2005, "A Simulation-Based Approach to Forecasting the Next Great San Francisco Earthquake", *Proceedings of the National Academy of Sciences*, 102(43), pp. 15363–15367.
- Sakoui, S.**, 2008, "Cat Bond Revamp for Blue Coast", *Financial Times* (July 29).
- Schwartz, J. M. and T. O. Wood**, 2008, "Life Settlements: Pricing Challenges and Opportunities", *The Journal of Structured Finance*, 14(2), pp. 70–76.
- Scott, J. S., J. G. Watson and W.-Y. Hu**, 2007, "Efficient Annuitization: Optimal Strategies for Hedging Mortality Risk", *Pension Research Council Working Paper No. PRC 2007–09*.
- Seitel, C. L.**, 2006, "Inside the Life Settlement Industry: An Institutional Investor's Perspective", *The Journal of Structured Finance*, 12(2), pp. 38–40.
- Seitel, C. L.**, 2009, "A Provider's Reflection from Inside the Life Settlement Industry: Understanding the Chaotic Environment", *The Journal of Structured Finance*, 15(2), pp. 70–73.
- Sharpe, W. F.**, 1964, "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk", *Journal of Finance*, 19(3), pp. 425–442.
- Sharpe, W. F.**, 1999, *Portfolio Theory and Capital Markets*, (New York: McGraw-Hill).
- Skypala, P.**, 2008, "Pensions Put Longevity Near Top of Hazards List", *Financial Times* (November 16).
- Standard & Poor's**, 2005, "New Issue: Box Hill Life Finance PLC", Standard & Poor's (January 31).
- Standard & Poor's**, 2006, "Ratings Raised On Vita Capital II's Catastrophe-Indexed Notes", Standard & Poor's (May 12).
- Standard & Poor's**, 2007, "Framework for Rating Natural Peril Catastrophe Bonds", Standard & Poor's (July 5).
- Standard & Poor's**, 2008, "Default Table Used to Rate Insurance-Linked Securitizations, Updated", Standard & Poor's (May 8).
- Standard & Poor's**, 2008, "Methodology and Assumptions for Rating Natural Catastrophe Bonds", Standard & Poor's (May 12).

- Standard & Poor's**, 2008, "Criteria Revised For Single-Event Natural Peril Catastrophe Bonds And Sidecar Transactions", Standard & Poor's (May 14).
- Standard & Poor's**, 2008, "Guide to Rating Process for Insurance-Linked Securities", Standard & Poor's (September 5).
- Standard & Poor's**, 2008, "Methodology and Assumptions Used for Rating Natural Catastrophe Insurance-Linked Securities", Standard & Poor's (September 11).
- Standard & Poor's**, 2008, "Guide To Rating Insurance-Linked Mortality Catastrophe Bonds", Standard & Poor's (September 11).
- Standard & Poor's**, 2008, "Insurance-Linked Securities – Capital Treatment and Basis Risk Analysis", Standard & Poor's (September 12).
- Standard & Poor's**, 2008, "Approach to Rating Indemnified Natural Catastrophe Insurance-Linked Securities", Standard & Poor's (September 15).
- Standard & Poor's**, 2008, "Clarifying The Framework For Rating Natural Catastrophe Bonds", Standard & Poor's (November 26).
- Stone, C. and A. Zissu**, 2006, "Securitization of Senior Life Settlements: Managing Extension Risk", *The Journal of Derivatives*, 13(3), pp. 66–80.
- Stone, C. and A. Zissu**, 2007, "The Return on a Pool of Senior Life Settlements", *Journal of Structured Finance*, 13(2), pp. 62–70.
- Stone, C. A.**, 2009, "The Supply and Demand for Life Settlement Contracts", *The Journal of Structured Finance*, Summer, 15(2), pp. 101–111.
- Stummer, C. and M. Sun**, 2005, "New Multiobjective Metaheuristic Solution Procedures for Capital Investment Planning", *Journal of Heuristics*, 11, pp. 183–199.
- Sullivan, R.**, 2009, "UK's First Longevity Swap Draws a Crowd", *Financial Times* (May 24).
- Swiss Re**, 2006, "Securitization: New opportunities for Insurers and Investors", *sigma*, 7/2006, Swiss Re.
- Swiss Re**, 2007, "Natural Catastrophes and Man-Made Disasters in 2006: Low Insured Losses", *sigma*, 2/2007, Swiss Re.
- Swiss Re**, 2009, "Scenario Analysis in Insurance", *sigma*, 1/2009, Swiss Re.
- Swiss Re**, 2009, "The Role of Indices in Transferring Insurance Risks to the Capital Markets", *sigma*, 4/2009, Swiss Re.
- Taleb, N. N.**, 2007, *The Black Swan: The Impact of the Highly Improbable*, (New York: Random House).
- Taleb, N. N.**, 2008, *Fooled by Randomness: The Hidden Role of Chance in the Markets and in Life*, (New York: Random House).
- Treaster, J.**, 2002, "Rethinking Dire Warnings by Insurers After September 11", *New York Times* (February 27).
- Uryasev, S.**, 2000, "Conditional Value-at-Risk: Optimization Algorithms and Applications", *Financial Engineering News Magazine*, 14 (February), pp. 1–7.
- Vadiveloo, J.**, 2005, *The Life Settlements Market: An Actuarial Perspective on Consumer Economic Value*, (Hartford, CT: Deloitte Consulting LLP and University of Connecticut).
- Valdez, E., L. Wang and J. Piggott**, 2008, "Securitisation of Longevity Risk in Reverse Mortgages", *North American Actuarial Journal*, 12(4), pp. 345–371.
- Vaugirard, V. E.**, 2003, "Pricing Catastrophe Bonds by an Arbitrage Approach", *The Quarterly Review of Economics and Finance*, 43(1), pp. 119–132.

- Venter, G.**, 2003, "Quantifying Correlated Reinsurance Exposures with Copulas", *Casualty Actuarial Society Forum*, Spring, pp. 215–229.
- Verbout, S. M., Schultz, D. M., Leslie, L. M., Brooks, H. E., Karoly, D. J. and K. L. Elmore**, 2007, "Tornado Outbreaks Associated with Landfalling Hurricanes in the North Atlantic Basin: 1954–2004", *Journal of Meteorology and Atmospheric Physics*, 97(1–4), pp. 255–271.
- Wang, C., and S. K. Lee**, 2008, "Global Warming and United States Landfalling Hurricanes", *Geophysical Research Letters*, 35(1), L02708.
- Wang, C. P., D. Shyu, Y. C. Liao, M.-C. Chen and M. L. Chen**, 2003, "A Model of Optimal Dynamic Asset Allocation in a Value-at-Risk Framework", *International Journal of Risk Assessment and Management*, 4(4), pp. 301–309.
- Wang, S.**, 1996, "Premium Calculation by Transforming the Layer Premium Density", *ASTIN Bulletin*, 26(1), pp. 71–92.
- Wang, S.**, 1998, "Aggregation of Correlated Risk Portfolios: Models and Algorithms", *Proceedings of the Casualty Actuarial Society*, LXXXV, pp. 848–939.
- Wang, S.**, 2000, "A Class of Distortion Operations for Pricing Financial and Insurance Risks", *The Journal of Risk and Insurance*, 67, pp. 15–36.
- Wang, S.**, 2002, "A Universal Framework for Pricing Financial and Insurance Risks", *ASTIN Bulletin*, 32, pp. 213–234.
- Wang, S.**, 2004, "Cat Bond Pricing Using Probability Transforms", *The Geneva Papers on Risk and Insurance*, Special Issue on "Insurance and the State of the Art in Cat Bond Pricing", 278, pp. 19–29.
- Wang, S.-M., J.-C. Chen, H.-M. Wee and K.-J. Wang**, 2006, "Non-linear Stochastic Optimization Using Genetic Algorithm for Portfolio Selection", *International Journal of Operations Research*, 3(1), pp. 16–22.
- Washington, S.**, 2007, "The Concept of a Model Act: The NAIC's Amended Viatical Settlements Model Act", *Journal of Structured Finance*, 13(2), pp. 89–92.
- Wighton, D.**, 2006, "The Holy Grail of the Asset Classes", *Financial Times* (November 29).
- World Economic Forum**, 2008, *Convergence of Insurance and Capital Markets*, World Economic Forum.
- World Economic Forum**, 2008, *Global Risks 2008: A Global Risk Network Report*, World Economic Forum.
- Woo, G.**, 1999, *The Mathematics of Natural Catastrophes*, (London: Imperial College Press).
- Worthington, A. and A. Valadkhani**, 2004, "Measuring the Impact of Natural Disasters on Capital Markets: An Empirical Application Using Intervention Analysis", *Applied Economics*, 36(19), pp. 2177–86.
- Yago, G. and P. Reiter**, 2008, "Financial Innovations for Catastrophe Risk: Cat Bonds and Beyond", *Milken Institute Financial Innovations Lab Report*, 5 (April), Santa Monica, CA.
- Yang, X.**, 2006, "Improving Portfolio Efficiency: A Genetic Algorithm Approach", *Computational Economics*, 28, pp. 1–14.
- Zaffino, P.**, 2009, "Catastrophe Bonds - Safe From The Storm", Guy Carpenter (March 24).
- Zeller, W.**, 2008, "Securitization and Insurance – Characteristics of Hannover Re's Approach", *The Geneva Papers on Risk and Insurance*, 33(1), pp. 7–11.
- Zeng, L.**, 2000, "Pricing Weather Derivatives", *The Journal of Risk Finance*, 1(3), pp. 72–78.
- Zeng, L.**, 2000, "On the Basis Risk of Industry Loss Warranties", *The Journal of Risk Finance*, 1(4), pp. 27–32.

Ziser, B., 2005, "Life Settlements: An Option for Seniors, an Opportunity for Investors", *The Journal of Structured Finance*, 11(2), pp. 12–14.

Zolkos, R., 2006, "Mad Scramble for Capital Fuels Cat Bond Market", *Business Insurance* (June 16).

Index

(page numbers in italic face refer to figures, panels and tables)

9/11 101, 238

A

Actuarial Guideline 38, *see* AXXX

AIDS and HIV 238, 247

Atlantic Multidecadal Oscillation
82–3, 84

AXXX reserves 200, 210–11

B

Bath's Law 60, 68

BNP Paribas EIB 366–70, 367

Buffett, Warren 7

C

capital strain 203–4

catastrophe bonds 9, 23–54; *see also*

catastrophe derivatives;

catastrophe risk

analysis of 415

credit risk in 167–79

collateralised reinsurance and,
collateral options in 177

customised puttable notes and
176

ILS and 169

Libor-linked returns and 177

new solutions needed for
170–1

solutions to, in insurance-

linked securities 171–6

traditional solutions to 169–70

Treasury money market funds
as collateral 176–7

trends and expectations 178–9

triparty repo arrangement and
173–6, 174

ways of mitigating 167–8

idea behind 24

investment performance of 42–6

issued by USAA 26

remodelling of 109–10

risk analysis of, illustrative

summary output of 107

structure of 28–31

catastrophe derivatives:

basis risk and 147–8

credit risk and 147

defined 120–1

index and, role of 120

index-linked 119–20

industry-loss warranties (ILWs)
and 119–52

defined 122

investor and hedger perspectives
on 150

investor universe and 149

key indexes and 123–5

- CME 127
- mortality and longevity 127–9
- PERILS AG 126–7, 126
- Property Claim Services (PCS)
 - 122, 123–5, 125
- Swiss Re and Munich Re 127
- market size and 122–3
- mortality and longevity
 - derivatives and 149–50
- pricing of, comments on 146
- reinsurance versus 121–2
- transformers and 148–9
- trends and expectations 150–2
- unusual products 145–6
- catastrophe risk:
 - bond remodelling 109–10
 - bond structure 28–31
 - bonds, historical performance of
 - 42–3, 43
 - bonds: trends and expectations
 - 52–4
 - default triggers 31–4
 - choice of 34
 - indemnity 31–2
 - index 32–4
 - perils and, number and types
 - of 34–5
 - second- or third-event 34
 - difference of, from other
 - investments 397
 - excess spread 44
 - historical perspective 25–6
 - managing portfolios of 391–433
 - additional considerations in
 - 429–30
 - beta, alpha and the Sharpe Ratio and 395
 - cat risk by (re)insurance company 404–7
 - constraints on 411–13
 - cumulative probability
 - distribution 423
 - decision-making process 424
 - diversification and variance of
 - portfolio return and 394
 - exotic beta and 396–7
 - individual securities and
 - 413–16
 - insurance-linked securities
 - 407–9
 - linear programming and
 - 418–19, 419
 - Markowitz efficient frontier and 393
 - misleading track records in
 - 431–2
 - performance measurement of
 - 431–2
 - portfolio construction 391–5
 - portfolio optimisation
 - 416–26
 - portfolio optimisation, pitfalls
 - of standard techniques in
 - 426–7
 - portfolio optimisation,
 - remodelling and 427–8
 - risk–return parameters 421
 - scenario testing 428–9
 - sensitivity analysis 428–9
 - standard tools for 413–16, 418
 - standard tools for, risks not
 - modelled using 416
 - statistics, examples of 420–2
 - types of instruments for
 - 409–10
 - VaR and TVaR for 422
 - market stability and growth
 - 46–7

-
- measures of, coherent 402–3
 - defined 403
 - model structure 97–8, 99
 - modelling 51–2, 55–118, 58
 - challenge of 55–6
 - cyclones and hurricanes 88–94
 - damage (hurricanes) 94–5
 - data quality and 106–8
 - earthquake frequency and severity and 59–61, 60
 - earthquakes 65–9
 - evolution of investor views on 86–8
 - financial loss (hurricanes) 95–7
 - hurricanes and cyclones 88–94
 - importance of, to investors 56–7
 - industry losses 129–30, 129
 - insurance-linked securities and 57
 - investor and 108–9
 - practical 103–8
 - as source of competitive advantage to investors 113–16
 - sponsor perspective on 112–13
 - technology 51
 - trends and expectations 116–18
 - practical model of 103–8
 - data quality and 106–8
 - property, bonds 23–54, 29, 30, 36
 - investment performance of 42–6
 - term 35
 - trends and expectations 52–4
 - Wang transform and pricing of 40–1
 - quantitative analysis and 35–44, 37, 38
 - exceedance curve 37–41
 - return period 41
 - stress testing and sensitivity analysis 41–2
 - by (re)insurance company 404–7
 - professional underwriters for 406–7
 - return and, comments on 403–4
 - return and, measures of 398–404
 - science of catastrophes and 57–9
 - sponsor and investor
 - perspectives and 47–50
 - diversification 47–8
 - slicing and packaging 48–50
 - types of 50
 - time-bomb, measures of 401–2
 - transfer of, in insurance 26–8
 - transfer of, in reinsurance 27–8, 27, 28
 - transferring to capital markets 24
 - insurance company motivation for 24
 - investor motivation for 24–5
 - volatility-related 399–401
 - see also* earthquakes; hurricanes and cyclones; property risk; terrorism; tsunamis
 - climate change 111–12
 - CME hurricane derivatives 140–4
 - contract types 141–2
 - geographical regions 142–3
 - other considerations involving 143
 - CME hurricane index 127, 128
 - collateralised reinsurance and, collateral options 177
-

- collateralised reinsurance, options in 177
- credit risk:
 - in catastrophe bonds 167–79
 - collateralised reinsurance and, collateral options in 177
 - customised puttable notes and 176
 - ILS and 169
 - Libor-linked returns and 177
 - new solutions needed for 170–1
 - solutions to, in insurance-linked securities 171–6
 - traditional solutions to 169–70
 - Treasury money market funds as collateral 176–7
 - trends and expectations 178–9
 - triparty repo arrangement and 173–6, 174
 - ways of mitigating 167–8
 - in insurance-linked securities 167–79
- Credit Suisse index 373–4
- customised puttable notes and 176
- cyclones and hurricanes 70–5, 78
 - forecasting 110–11
 - frequency and severity effects of, over various time horizons
 - 81–3
 - medium term 82–3
 - medium–long term 83
 - short term 81–2
 - historical frequency of (US) 74–7
 - Katrina 19, 25, 26, 34, 46, 46–7, 48, 84, 124, 130, 131, 160, 164, 165, 238, 400, 432
 - landfall of, in peak regions 79–80
 - macro-scale severity and effects of, investor views on 83–5
 - modelling of 88–94
 - damage 94–7
 - financial loss 95
 - frequency 89–90
 - frequency and intraseasonal correlation 90–1
 - wind field 91–4, 92
 - North Atlantic distribution of 79
 - Rita 130, 160, 164
 - SE US return periods of 80
 - seasonality of risk of, in insurance-linked securities 77–9
 - Wilma 19, 130, 160, 164
- D**
 - damage modelling (hurricanes) 94–5
 - default triggers 31–4
 - Deutsche Börse Xpect index 375–6
 - diversification and variance of portfolio return 394
- E**
 - earthquakes:
 - frequency and severity of 59–61, 60
 - location of 61–5
 - modelling of 65–9
 - scales of, examples 61
 - simulating 67, 70
 - tectonic plates and 61, 62, 64
 - tsunamis caused by 69–70
 - wind and, structural engineering analysis of 96
 - El Niño 81–2, 82, 84, 85

-
- embedded value (EV)
 - securitisation 213–33
 - closed block and 216
 - defined 214–15, 215
 - examples of 227–9
 - Gracechurch/Barclays 227–9, 229
 - Scottish Equitable 230–2, 231
 - investor perspective on 216–17
 - modelling 220–3
 - assumptions for 222–3
 - cashflow 223–4
 - non-actuarial risks 224
 - ratings of 225–7
 - caps 226–7
 - surveillance 227
 - rationale for 231
 - stress scenarios 224–5
 - structures of 217–20, 219, 221
 - trends and expectations 232
 - true, versus direct monetisation 215–6
 - value-in-force (VIF) and 214–15, 215
 - embedded-value and value-in-force securitisation 202–3
 - EUREX hurricane futures 144
 - EV, *see* embedded value
 - excess insurance reserves 199–212
 - AXXX reserves and 200, 210–11
 - capital strain and 203–4
 - embedded-value and value-in-force securitisation 202–3
 - examples of 199–201
 - “excess” 201
 - funding solutions 201–2
 - investors in, additional
 - considerations for 209–10
 - loss portfolio transfer and 211
 - market fluidity and 203
 - RBC requirements and 203–4
 - Regulation XXX and 204–8, 205
 - letters-of-credit facility for funding 206–7
 - securitisation of reserves 207–9, 208
 - exotic beta 396–7
 - extreme mortality risk:
 - basis risk and 246–7
 - bond structure 239
 - credit enhancement and 247–8
 - current modelling approaches to 251–9, 252
 - baseline 252–3, 252
 - pandemic 254–7, 255, 256, 257
 - terrorism 253–4, 253, 254
 - explained 237–8
 - factors affecting 249–50
 - growing awareness of 238
 - HIV / AIDS and 238, 247
 - investor types 248
 - investors in, considerations for 259–60
 - modelling results, analysis of 258
 - mortality derivatives and 259
 - other securitisations of 244–6, 244
 - probabilistic modelling of, in
 - securitisation context 258
 - quantification and pricing of 248–50
 - rates modelling and 249
 - reference index construction for 243
 - scenario testing 258–9
 - securitisation of 237–61
 - basis risk and 246–7
 - Vita Capital transaction 240–2, 240, 244
-

- Tartan Capital transaction and 245, 245, 246
- trends and expectations 260
- F**
- financial-loss modelling
 - (hurricanes) 95–7
 - aggregate approach 97
 - demand surge and 96–7
- G**
- Gracechurch/Barclays 227–9, 229
- Gray, William 110
- Gutenberg–Richter law 59, 60, 60, 63, 68
- H**
- heating and cooling degree days
 - (HDD and CDD) 183–7, 186, 187
- HIV/AIDS and 238, 247
- Hurricane Katrina 19, 25, 26, 34, 46, 46–7, 48, 84, 124, 130, 131, 160, 164, 165, 238, 400, 432
- Hurricane Rita 130, 160, 164
- Hurricane Wilma 19, 130, 160, 164
- hurricanes and cyclones 70–5, 78
 - forecasting 110–11
 - frequency and severity effects of, over various time horizons
 - 81–3
 - medium term 82–3
 - medium–long term 83
 - short term 81–2
 - historical frequency of (US) 74–7
 - Katrina 19, 25, 26, 34, 46, 46–7, 48, 84, 124, 130, 131, 160, 164, 165, 238, 400, 432
 - landfall of, in peak regions 79–80
 - macro-scale severity and effects of, investor views on 83–5
 - modelling of 88–94
 - damage 94–7
 - financial loss 95
 - frequency 89–90
 - frequency and intraseasonal correlation 90–1
 - wind field 91–4, 92
 - North Atlantic distribution of 79
 - Rita 130, 160, 164
 - SE US return periods of 80
 - seasonality of risk of, in insurance-linked securities 77–9
 - Wilma 19, 130, 160, 164
- I**
- index-linked catastrophe
 - derivatives 119–20
- industry-loss warranties (ILWs):
 - catastrophe derivatives and 119–52
 - defined 122
 - market for 130–2
 - structuring 130–2
 - US windstorm and 131
- insurable interest 275–7
- Insurance Futures Exchange (IFEX)
 - 133–40 *passim*, 137, 138, 139, 144, 151
- block trades 140
- catastrophe derivatives 133–40
- contract specifications 134–7, 135–6
- event-linked futures,
 - maintenance margin requirements 139

-
- margin requirements 137
 - US windstorm settlement prices 137
 - insurance risk, no clear definition of 4
 - insurance risk, transfer of, to capital markets, reasons for 14–17
 - insurance markets:
 - explained 4–6
 - securities issued within 6–7
 - insurance-linked securities (ILS) 7–9, 8, 13–20, 15
 - broader definition of 181
 - catastrophe, managing portfolio of 407–9
 - correlation among different types of 438
 - credit risk in 167–79
 - defined 13
 - different types of, in same portfolio, rationale for 437–8
 - flexibility and diversification 437
 - skill set and expertise 437–8
 - diversification and yield offered by 17–19
 - examples of 9
 - hurricanes and, seasonality of risk of 77–9
 - investment management policy and 446
 - investor sophistication and analytical expertise 87
 - market dynamics and 19–20
 - market development, drivers of 450–1
 - market growth, obstacles to 451
 - market, state and future development of 450
 - mortality and longevity models
 - in 295–322, 306, 312
 - age transform 319–20
 - credibility theory approach 310–12
 - dynamics 306–8, 307, 308
 - Lee–Carter method and 315–16
 - longevity improvements 312–15, 314
 - Markov process and 316–19
 - models, distributions, tables, basic concepts of 296–9
 - population to individual, moving from 318–19
 - rates 296, 299
 - shocks 320–1
 - tables 299–301, 300–1, 311
 - tables, population 301–5, 303, 304–5
 - tables, select and ultimate 308–10
 - multiple types of, in same portfolio 441
 - multiple types, portfolios of 435–48
 - optimisation of 439–41
 - opportunities in 451
 - as portfolio diversifier 11
 - portfolio monitoring, stress testing and 447
 - portfolio performance
 - measurement 442–5, 444, 445
 - standard types of 443–4
 - portfolio valuation issues concerning 441–2
-

- risk management and 446–7
 - sensitivity analysis and 440
 - tenor and liquidity of 438–9
 - types of 14–17, 435–7, 436
 - valuation issues concerning 441–2
 - various types of, correlation among 438
 - yield and diversification offered by 17–19
 - efficient frontier 18–19
 - yield generation 18
 - zero-beta assets 18
- International Swaps and Derivatives Association (ISDA) 132–3
- wind swap confirmation template of 132
- K**
- Katrina, *see* Hurricane Katrina
- L**
- La Niña 81–2
- Lee–Carter method 315–16
- legal and ethical issues concerning
 - life insurance 267–9
- letters-of-credit facility for Regulation XXX reserves 206–7
- life insurance:
 - contestability and 280–1
 - evolution of market in 265
 - industry perspective 289–90
 - insurable interest and 275–7
 - investor due diligence and 281–2
 - investor- or stranger originated 278–9
 - premium financing and the STOLI issue 278–80
 - wet paper and 279–80
- investor perspective 285–7
 - competitive advantages and disadvantages 288–9
 - historical investment performance 286
 - overstated portfolio values 287–8
 - portfolio valuation 286–7
- investor risk/consumer protection and, link between 273–4
- legal and ethical issues concerning 267–9
- market participants 269, 270
- market size of, current and future 269, 271–2
- not-for-profit organisations and 282–5
- policy as tradable asset 16, 263–4
- regulatory issues concerning 272
- risks to insurers 290–2
 - expenses 292–3
 - investment assumptions 292
 - mortality and lapses 291–2
- securitisations 266
- settlements 263–94, 271
 - assumed premiums and 342
 - debits 328–30
 - knockout versus debit/credit approach to 328
 - life expectancy (LE) and 325–7
 - life expectancy (LE) and, choosing 340–2
 - methodology changes in calculation of 327–8

-
- modelling investment
 - performance of 323–5
 - mortality table and, choice of
 - 330–1
 - older ages and, underwriting
 - for 335–9
 - projection pursuit regression
 - 332
 - relative risk ratios (RRRs) and
 - 334–5
 - risk and 343
 - underwriting concepts
 - concerning 328
 - valuation basic table (2008)
 - 331–4
 - valuation of, and other
 - mortality-linked securities
 - 323–45
 - viatical settlements versus
 - 265–6, 266
 - Whitaker–Henderson method
 - and 334
 - tax issues concerning 274–5
 - trust structures and 281–2
 - life settlements 263–94, 271
 - assumed premiums and 342
 - debits 328–30
 - knockout versus debit/credit
 - approach to 328
 - life expectancy (LE) and
 - 325–7
 - additional bias 341–2
 - choosing 340–2
 - determination of 32–4
 - longevity derivatives in 381–3
 - managing longevity risk in
 - 380–1
 - methodology changes in
 - calculation of 327–8
 - modelling investment
 - performance of 323–5
 - mortality table and, choice of
 - 330–1
 - older ages and, underwriting for
 - 335–9
 - underwriting tools 337–9
 - additivity of debits and credits
 - 339
 - projection pursuit regression 332
 - relative risk ratios (RRRs) and
 - 334–5
 - risk and 343
 - securitisation of 383
 - underwriting concepts
 - concerning 328
 - valuation basic table (2008)
 - 331–4
 - valuation of, and other
 - mortality-linked securities
 - 323–45
 - viatical settlements versus 265–6,
 - 266
 - Whitaker–Henderson method
 - and 334
 - LifeMetrics index 374–5
 - longevity:
 - bonds 364–71, 368
 - BNP Paribas EIB 366–70, 367
 - pricing, comments on 370–1
 - hedges, natural, when writing
 - contracts for 355–6
 - hedges, standardised index-
 - based 364
 - improvements 312–15, 314,
 - 352–4, 352
 - modelling 353–4
 - indexes of 373–6
 - Credit Suisse 373–4
-

- Deutsche Börse Xpect 375–6
 - LifeMetrics 374–5
 - investors in 376–8
 - life settlements and 380–3
 - market developments
 - concerning 378–9
 - models of, in insurance-linked
 - securities 295–322, 306, 312
 - mortality forwards, survivor
 - forwards and 359–64, 360, 361, 362, 363
 - rates, term structure of 363–4
 - risk 347–52
 - definition of 347–8
 - entities and securities exposed to 348
 - need to transfer 349–52
 - risk transfer, primary
 - mechanisms of 355–7, 356
 - risk transfer and securities
 - linked to 347–87
 - solutions to risk management of,
 - in a DB pension plan 371–3
 - swaps 357–9
 - cashflow exchange in 358
 - traded policies and, extension
 - risk in 379–83
 - trends and expectations 384–7
- M**
- market dynamics 19–20
- market stability and growth 46–7
- Markov process 316–19
- Markowitz efficient frontier 393
- mortality and longevity derivatives
 - 149–50
- mortality and longevity models
 - 295–322, 306, 312
 - age transform 319–20
 - credibility theory approach 310–12
 - dynamics 306–8, 307, 308
 - Lee–Carter method and 315–16
 - longevity improvements 312–15, 314
 - Markov process and 316–19
 - mortality modelling using physiological age 316–18, 317
 - models, distributions, tables,
 - basic concepts of 296–9
 - population to individual, moving from 318–19
 - rates 296, 299
 - shocks 320–1
 - tables 299–301, 300–1, 311
 - tables, population 301–5, 303, 304–5
 - tables, select and ultimate 308–10
- mortality derivatives 259
- O**
- Omori–Utsu law 59, 61, 68
- P**
- pandemic flu 101–3
- PERILS AG 126–7, 126
- portfolio diversification through
 - insurance-linked securities 11
- Property Claim Services (PCS) 122, 123–5, 125
- property risk:
 - catastrophe bonds concerning 23–54; *see also* catastrophe risk

-
- investment performance of
 - 42–6
 - structure of 28–31
 - sponsor and investor
 - perspectives and 47–50
 - diversification 47–8
 - slicing and packaging 48–50
 - types of 50
 - market stability and growth
 - 46–7
 - securitisation of 23–4
 - R**
 - reference index construction
 - 243
 - Regulation XXX 170, 200
 - reserve funding 204–7, 205, 206
 - letters-of-credit facility for
 - 206–7
 - securitisation of 207–9, 208
 - reinsurance:
 - investment analysis and,
 - considerations in 162–3
 - investor universe and 160–2
 - securitised, reinsurance
 - sidecars and 153–66, 161
 - investor perspective 156–7
 - leveraged versus equity-only
 - 159–60
 - representative transactions in
 - 160
 - sponsor perspective 157–8
 - structure of 154–6, 155
 - types of 158–9
 - trends and expectations and
 - 163–6
 - relative risk ratios (RRRs) 334–5
 - risk as good thing 3–4
 - Rita, *see* Hurricane Rita
 - S**
 - Saffir–Simpson Hurricane Wind
 - Scale 71, 72–3, 73, 74, 75, 76, 79, 128
 - San Andreas fault 60, 64, 66, 114
 - securities from insurance
 - companies 6–7
 - securitisation of extreme mortality
 - risk 237–61
 - basis risk and 246–7
 - benefits of, to Swiss Re 242–3
 - distribution by age within
 - 241
 - geographic distribution within
 - 242
 - payout schedule for 242
 - structure of 240–1, 240
 - trigger index for 241
 - Vita Capital transaction 240–2, 240, 244
 - securitised reinsurance:
 - investment analysis and,
 - considerations in 162–3
 - investor universe and 160–2
 - reinsurance sidecars and 153–66, 161
 - investor perspective 156
 - leveraged versus equity-only
 - 159–60
 - representative transactions in
 - 160
 - sponsor perspective 157–8
 - structure of 154–6, 155
 - types of 158–9
 - trends and expectations 163–6
 - seismic hazards map 63
 - September 11 attacks 101, 238
 - stranger-originated life insurance (STOLI) 273
-

Swiss Re, Vita Capital transaction
of 240–2

T

Tartan Capital 102
mortality rate definition for 245
tectonic plates 61, 62, 64
terrorism risk 98–101, 253–4
transfer of insurance risk to capital
markets, reasons for 14–17
Treasury money market funds as
collateral 176–7
tsunamis 69–70, 69

V

Vita Capital transaction 240–2, 240,
244

W

weather derivatives 181–96
defined 181–3
emissions trading and 193–4
exchange-traded 187–8
heating and cooling degree days
(HDD and CDD) 183–7, 186,
187

explained 183
investing in 191–3
specific strategies for 192
valuation and 192–3
other types of 184–5
pricing models for 188–9
burn analysis 188–9
stochastic temperature
189
pricing, practical challenges in
189–90
data issues 190
forecasts 191
time period of historical
observations, choice of
190–1
standards options, payout on
186
trends and expectations 194–6
types of underlyings 182–3
weather insurance and 182
World Food Programme/AXA
Precipitation 185
Wilma, *see* Hurricane Wilma
World Trade Center attack (9/11)
101, 238