AJAP

AUSTRALIAN JOURNAL OF ACTUARIAL PRACTICE | VOLUME 1

Actuaries Institute

The Australian Journal of Actuarial Practice (AJAP)

The AJAP is the journal of the Actuaries Institute and is aimed at leading debate in areas where actuaries practise in Australia so as to enhance the work of practitioners and improve the service provided to their employers, their clients and the community. The AJAP will publish papers, notes and commentary. All content of the AJAP is subject to a peer review process.

The AJAP will publish electronic volumes at least twice annually.

The Australian Journal of Actuarial Practice is published by
The Institute of Actuaries of Australia
Level 7 Challis House, 4 Martin Place Sydney NSW 2000 Australia
T +61 (0) 2 9233 3466
W www.actuaries.asn.au
E AJAP@actuaries.asn.au

Design
Kirk Palmer Design

ISSN 2203-5354

Editorial Policy

The following types of articles will be considered for publication:

- **Papers** of generally between 2,000 and 5,000 words that deal with some aspect of professional practice relevant to members of the Actuaries Institute and where the author(s) bring insights into how this aspect could be improved or better interpreted. Papers may focus on new applications of actuarial skills or on improving actuarial practice in more traditional areas, on the interpretation and implementation of substantial regulatory issues relevant to members and on the application of newly developed theory, methods or techniques from other disciplines to actuarial work.
- **Notes** of generally between 250 and 1,000 words that comment on matters of interest to members of the Actuaries Institute with the intention of bringing to members a new idea or understanding of a topic.
- **Discussions on published papers and notes of between 100 and 500 words**. Discussions can be emailed to AJAP@actuaries.asn.au

All submissions to the AJAP are peer reviewed.

Full guidelines for authors are available from the Actuaries Institute.

Important Information for Contributors

The Editorial Committee reserves the right to accept, reject or request changes to all submissions as well as edit articles for length, basic syntax, grammar, spelling and punctuation.

AJAP 2014 VOLUME 1

CONTENTS

Go to **www.actuaries.asn.au/ knowledge-bank/journals** and hear Editor John Evans' welcome remarks.

FROM THE EDITOR

Welcome to the first edition of the *Australian Journal of Actuarial Practice*, which succeeds the *Australian Actuarial Journal*.

The AJAP is a scholarly journal containing original papers, shorter notes and comments on previous papers with an emphasis on issues of concern to practising actuaries. Each edition of the AJAP is managed by a subcommittee of the Editorial Committee, and on this occasion I would like to thank my colleagues Rade Musulin and Timothy Kyng for their very significant efforts in reviewing papers and assisting with the publication.

I would also like to thank the Institute team, Katrina McFadyen and Nicole Sitosta for the great assistance in the organisation and production of the AJAP.

In this edition, primarily focused on 'risk', we have papers relating to a wide area of risk issues, with an innovative paper looking at how to measure risk using microsimulation data, a paper of significant interest to actuaries in Australia on how to extrapolate the yield curve, and two papers relating to risks involving hedging, with notes relating to a survey carried out on how insurers regard human capital risk and stress testing for insurers.

We have also invited notes on issues of broader interest to actuaries, relating to superannuation taxation, disability management and graduates' views of their actuarial education, with the intention this will generate comments for publication in further editions.

'Risk', quite rightly, has become a major focus of institutions and regulators, with the ultimate goal that risk management is fully integrated into the business management process. Undoubtedly the banks that were faced with the Basel regulations, commencing over a decade ago, are more advanced in this business integration, having learned that prudential regulation of risk was not sufficient for the risk management of their business. It is early yet, but there are signs that the Australian banks are developing their thinking about risk management and how it can better serve the business well beyond prudential requirements and are moving to consider alternative processes to the historically driven quantitative based approaches which may well have misled businesses into believing they had modelled all risks. Insurers are rapidly coming to grips with developing risk management approaches and its integration into the business process, but the development varies significantly across the industry. Superannuation funds are just starting the journey and have a long way to go in developing integrated business risk management systems that are owned and understood by the Board, management and employees.

The Note by Dr Royal highlights one of the significant issues now emerging in risk management as businesses attempt to develop systems that assist with management of their businesses and by necessity need to then be predictive. Dr Royal's survey clearly shows a lack of understanding by institutions of the very complex interactions taking place, and in that particular instance, the significant role that human capital management plays in exposure to risk events.

There is a need for both quantitative and qualitative analysis in the risk management process. The real value of the quantitative analysis is not its predictive ability, which is usually poor, but its ability to help businesses understand the drivers of their risk events, as these may not be clear in a complex institution, and certainly not in their relative importance. By understanding the drivers of risk events, qualitative analysis based on systems analysis can be used to understand the likely impact of the identified drivers in the dynamic world of a financial institution.

It is intended the establishment of the Journal will allow well thought out and peer reviewed discussion of issues pertaining to the work of actuaries to assist in the evolution of the thinking of actuaries on matters of importance to their employers and the community.

John Evans
EDITOR

Measuring risk using household microsimulation models

DR DAVID SERVICE, DR RICHARD CUMPSTON AND HUGH SARJEANT

Dr David Service

ABSTRACT

The more we know about a problem the more chance we have of designing effective solutions. Models are tools for using our existing and emerging knowledge, and take many forms.

In particular, microsimulation models enable the measurement of risks where the only robust approach is to use scenarios. Microsimulation allows scenarios to be modelled in significant detail, which enables reliable risk measurement and subsequent management.

Models with fine geographic subdivisions can be used for many purposes, for example the selection of locations for infrastructure such as age-care facilities, and the simulation of risks for mortgage insurers.

The limited geographic subdivisions of census unit record samples have led to the recent development of a method to create synthetic persons and dwellings from census tabulations. This method has been applied to Australian and New Zealand census data.

Census data have been supplemented by survey data, for example on income, expenditure, wealth, housing values, mortgages and superannuation. Models of 123 diseases have been used to impute and project disease conditions.

Five specific applications of household microsimulation are discussed.

KEYWORDS
microsimulation, risk measurement, household microsimulation models, disease models

AJAP 2014; 1: 5-11

INTRODUCTION

Bellis, Lyon, Klugman & Shepherd (2010) describe the actuarial control cycle as based on a simple problem-solving algorithm

- define the problem
- design the solution
- monitor the results.

Actuaries employed by insurers are using complex models of their companies and products (Smith 2010). But actuaries may not be making systematic use of the increasing amounts of data available from many sources about Australians and their households. These large amounts of data may be best modelled using microsimulation models, rather than multistate models.

Australian household microsimulation models are increasing in size and functionality. A model proposed in 2006 had 180,000 sample persons and one area (Percival 2007), while a model currently under development has up to 22 million persons and 2200 areas (Service, Cumpston & Sarjeant 2013).

We suggest that household microsimulation models can be used by governments, businesses and not-for-profits to help

- integrate data from many sources
- identify market opportunities
- measure and manage complex risks.

While microsimulation models have important applications in actuarial fields, their range of use spans a wide spectrum of problems.

RISK MANAGEMENT

Risk management is an emerging science in its own right, although actuaries have been actually doing it for several hundred years.

The process comprises three parts – risk identification; risk measurement; and risk management. It is in the second part – risk measurement – that the use of microsimulation models can make a major contribution.

While actuaries may use stochastic techniques to measure many of the risks which they work with, there remain a large set of risks where these mathematical tools do not work successfully, e.g. major shifts in immigration, dramatic changes in workers' compensation legislation, changes in birth rates. In such cases the best measurement tool is the scenario-testing approach. But some scenarios are very difficult to model robustly.

These are the situations where microsimulation models can provide robust and timely measurement of potential scenarios resulting from the changes being investigated. This provides an excellent tool for the measurement of the risk arising should the change occur. Without such measurement the management of such risks becomes a matter of guesswork and is very 'hit and miss'.

As microsimulation is agent-based, complex risk interactions can be simulated. A reviewer of this paper noted that changing the fee structure for a particular product may impact the highest value customers in a complementary product. Observing these interactions can help identify risks as well as measure them.

This paper explores some of the important features of microsimulation models and provides examples of the uses of such models in real situations.

A CASE FOR SCENARIO-BASED MODELS WITH EXAMPLES FROM DEMOGRAPHIC AND SOCIAL TREND SCENARIOS

To illustrate some of the events which do not lend themselves to stochastic risk measurement, this section gives examples of recent demographic and social trends which can lead to risks – both positive and negative – emerging.

Although interesting, none of these examples provide any detailed guidance about appropriate risk measurement. Scenario testing is needed in order to make any measurement of the impact of such events occurring. Detailed models are needed, incorporating the information available from a range of sources. These can be multistate models, which are cell based, or microsimulation models, which are record based. Either type of model would, for example, allow inequality issues to be analysed at regional levels.

Life expectancies at birth have increased from 1970–72 to 2005–07 by about 0.31 years per year for males, and 0.25 for females (Australian Government Actuary 2009). Oeppen & Vaupel (2002) provide international data on continuing linear increases in life expectancies, and comment that

> The officials responsible for making projections have recalcitrantly assumed that life expectancy will increase slowly and not much further. The official forecasts distort people's decisions about how much to save and when to retire. They give politicians licence to postpone painful adjustments to social-security and medical-care systems.

Total fertility rates were 2.89 in 1966, fell to 1.73 in 2001, and rose to 1.88 in 2011 (Australian Bureau of Statistics 2012d). As a result, there will be fewer children available to care for ageing parents.

Leigh (2013 p113) notes that 19% of children live in lone-parent households, up from about 10% in the late 1970s. He attributes this mostly to lower partnering rates, rather than the slight increase in divorce rates since no-fault divorce was introduced in 1975.

The share of total incomes received by the top 1% of Australian adults increased from about 4.5% in 1983 to 9% in 2010 (Leigh 2013 p35). The share of wealth held by the top 1% increased from about 7% in 1978 to 11% in 2010 (p55). Income distribution is, however, more equal than in the United States, where Stiglitz (2013 p20) estimates the top 1% receive about 20% of national income.

MULTISTATE MODELS

Multistate models are cell based, with each cell containing the number of persons with a particular combination of characteristics. Whitelock-Jones (2003) notes that, if transition probabilities depend on the duration spent in the cell, then the theoretical number of cells required in the model can become very large, and the model becomes very complex.

Large numbers of cells can also be needed to model relationships between individuals. Orcutt, Greenberger, Korbel & Rivlin (1961 p288) estimated that, allowing for combinations of race, marital status, parent ages, child ages, interval since marriage and female parity, a model of US families would have required at least a billion cells.

MICROSIMULATION MODELS

The International Microsimulation Association (2009) defines microsimulation as

> a modelling technique that operates at the level of individual units such as persons, households, vehicles or firms. Within the model each unit is represented by a record containing a unique identifier and a set of associated attributes – e.g. a list of persons with known age, sex, marital and employment status; or a list of vehicles with known origins, destination and operational characteristics. A set of rules (transition probabilities) are then applied to these units leading to simulated changes in state and behavior.

Two types of microsimulation are used – static and dynamic. Static microsimulations start with a set of records and model their short-term behaviour. For example, taxation and social security agencies use large samples of their client records to simulate the short-term effects of legislative changes. Changing the record weights can be used to allow for future client changes, for forecasts of up to about 5 years (O'Donoghue 2001 p9).

By contrast, dynamic microsimulations model changes to the unit records, by adding new records (for example births and immigrants) and removing some existing records (deaths and emigrants). Changes to individual characteristics such as education, employment and wealth can be modelled, as well as changes to dwelling characteristics and household occupants. Long-term projections are feasible.

COMPARISONS BETWEEN MULTISTATE AND MICROSIMULATION MODELS

The number of cells required by a multistate model is the product of the number of values modelled for each characteristic. For example, a model with 1000 areas, 110 ages, 2 sexes, 5 ethnic groups, 2 employment statuses and 13 income groups would require 1000*110*2*5*2*13 cells – about 29 million cells. Much larger numbers of cells may be needed for some purposes.

The number of data items stored for a microsimulation model is the number of records times the number of characteristics per record. For example, a model with 1000 areas might need about 2 million records to give at least 1000 records for most areas. The number of data items stored in the above example might thus be about 2 million by 6, or about 12 million.

The numbers of probability calculations in a multistate model will depend directly on the number of cells. For example, any individual has a probability of death, so that the number of persons in each cell would have to be recalculated to allow for deaths in each simulation cycle. By contrast, a microsimulation model would need a death probability calculation for each record.

Memory requirements and runtimes thus depend on the number of cells for multistate models, and on the number of records for microsimulation models. Cell numbers can be very large where characteristics have many possible values, or where relationships between individuals have to be modelled. Multistate models are used for population projections, sometimes with additional characteristics such as indigenous status (Wilson 2009). For many purposes, however, a microsimulation model will need less memory, and run more quickly.

While all scenario measurement relies on models, the fundamental question is the reliability of the measurement produced by the model employed. Microsimulation models provide the ability to deal with scenarios in detail and to produce reliable measurement as a basis for good risk management.

HOUSEHOLD MICOSIMULATION MODELS

A household microsimulation model starts with a record for each person in the baseline population, and a record for each dwelling. Each person record contains a link to the record of their dwelling, and may also contain kinship links, for example to parents, partners and children. Dwelling records may include vacant dwellings. During a projection, demographic events are randomly simulated – births, deaths, emigrants, immigrants, internal migrants, household entries and exits, movements to and from non-private dwellings. Education, employment, earnings, other income, expenditure, superannuation, housing purchases and other investments may be simulated. The creation of new dwellings, and the demolition of some existing dwellings, may be simulated. The output from a projection will be an updated set of person and dwelling records, together with any event totals of interest.

Household microsimulation models with fine geographic subdivisions

National models with more geographical subdivision or shorter simulation cycles can have a wide range of uses, provided their run times can be kept reasonable. For example

- Models with fine geographic subdivisions could be used to select viable locations for long-term facilities (such as schools, churches and age-care facilities).
- Models with fine geographic subdivisions and short time cycles could be used to simulate employment and housing markets.

Caldwell (1983 p61) commented

> Policy analysis payoffs become more substantial when spatial disaggregation reaches the state level in American models. But the gains accumulate exponentially when the level of spatial disaggregation becomes even finer than the state level.

He noted that

> The major difficulty with precise location identification is the need for far larger samples in order to reduce sample variance for area-specific outcomes. One might imagine working with one percent, five percent or even ten percent samples of the population.

Synthetic baseline populations for household microsimulation models

Confidentialised unit record files of persons and dwellings have been available for 1% samples from Australian censuses and have been used as baseline populations for household microsimulation models (Percival 2007; Cumpston 2011). These files have only had about 50 geographic regions, and for confidentiality reasons household sizes have been truncated, and some personal details changed or omitted. A 2% sample file was released for the 2001 New Zealand census, but not for the 2006 census. These limitations have recently led to the synthesis of large census sample files for Australia and New Zealand (Service, Cumpston & Sarjeant 2013).

The Australian Statistical Geography Standard (Australian Bureau of Statistics 2010), applying from 1/7/11, specifies about 237,000 mesh blocks, 55,000 statistical areas level 1 (SA1) and 2200 statistical areas level 2 (SA2). Census cross-tabulations are readily available for SA2s for the 2011 census, with a randomisation process being applied to each cell for confidentiality reasons. While some data are available for smaller areas, confidentialisation means that the SA2s are probably the smallest areas that can be used to synthesise populations.

With Australia's total population of about 22 million in 2011, a 10% sample size gives an average of about 1000 sample persons per SA2. Given the uneven size distribution of the SA2s, a 20% sample size may be needed to ensure that there are at least 1000 persons in most of the synthetic SA2s.

Adding data from other sources to baseline populations

The Australian Bureau of Statistics provides unit records in many confidentialised unit record files, including

- Australian Health Survey 2011–13 (2013a)
- Income and Housing 2011–12 (2013b)
- Household Expenditure Survey and Survey of Income and Housing 2009–10 (2012c)
- Disability, Ageing and Carers 2009 (2012b).

While only broad area information is provided in these unit record files, sufficient information about age, sex, income and family structure is usually provided to be able to reasonably impute values for each person in a synthetic baseline population.

Detailed aggregates with sex/age breakdowns for self-managed superannuation funds were available from the Australian Taxation Office (2011).

The HILDA Survey is a household-based longitudinal survey, named the Household Income and Labour Dynamics in Australia Survey. It has released unit record data for the first 11 waves (Summerfield et al 2012). Waves 2, 6 and 10 have included questions on wealth and superannuation. HILDA is very useful for imputing baseline values and as a source of transition assumptions.

Imputing and modelling diseases

The Australian Institute of Health and Welfare published its first burden of disease study in 1999 (Mathers, Vos & Stevenson). Released with this report were models for about 175 diseases, giving estimates for their incidence and duration, and some indication of the mortality rates for each disease. These models were not released for the second study, and we do not know whether they will be released for the current third study.

The disease models released by the first study were used by Cumpston (2011a) to impute and project 123 disease conditions, with 584 stages. Baseline diseases were imputed by simulating the incidence and development of diseases from birth for each person in the baseline file. Deaths were simulated from the mortality rates assumed for each person's diseases, rather than from assumed population mortality rates.

Synthesis and projection runtimes

Synthesis of a 1% sample of persons and households from the Australian 2011 census (about 220,000 persons located over 2200 SA2s) took about 30 seconds on a laptop computer. Projection for 50 years took about a minute. These trials included the imputation and projection of diseases. No use was made of parallel processing, which might have provided a 50% reduction in runtimes.

These runtimes, and the computer memory required, should increase linearly with the numbers of persons synthesised. The number of areas modelled should have little effect on runtimes or storage requirements. With enough computer memory, it should be feasible to synthesise and project 22 million Australians.

High development costs for household microsimulation models

Harding (2007 p4) noted that the developers of the MINT, CBOLT and POLISIM models of the US had estimated that the total development costs of each model had exceeded US$6m. She considered that budgets of this magnitude had also been involved for the Canadian DYNACAN model, the Norwegian MOSART model and the Swedish SESIM model. .The function of the DYNACAN model was to provide checks on actuarial estimates of Canadian social security costs, and it is no longer operational.

As noted above, fine geographic subdivisions can greatly increase the numbers of potential users, and thus the financial viability of the model. The marginal cost of adding finer geographic subdivisions is low.

Large numbers of data fields should also increase the number of potential users. Each field needs to be synthesised for the baseline population from census data, or imputed using other data. Assumed transition probabilities are needed to project each field. A new field, such as self-assessed indigenous status, may have complex effects on assumptions about fertility, health, mortality, education, employment and income. Marginal costs of adding fields may thus be high.

Validating and updating household microsimulation models

The data synthesis process automatically generates many tables and graphs, comparing aspects of the synthesised data with known data (Service, Cumpston & Sarjeant 2013). Similarly, the projection program generates many tables and graphs, comparing the year by year results with any available surveys or administrative data. To reduce random variations, comparisons may be made using the averages of multiple synthesis and projection runs.

Inevitably, projections will differ from prior expectations. For example, immigrant numbers, child births and superannuation contributions may all be strongly affected by legislative changes. The projected persons and dwellings need to be adjusted year by year to ensure that they correspond closely with current data. This is sometimes done by alignment, where the simulations for each type of event each year are adjusted to match emerging data. An alternative is to adjust one of the parameters for each process each year to give closely matching results. As new census and survey data become available, the baseline population should be resynthesised.

As in macro-level stochastic simulations, microsimulations are subject to parameter uncertainty. A range of techniques has been used to reduce runtimes, allowing many repeats of the same simulation. For example, stochastic mortality assumptions can be included. In the authors' work the data are synthesised using stochastic techniques, so different baseline data can be created for each run. These approaches enable the sensitivity of the results to parameter variations to be tested.

Improvements in data availability and household microsimulation techniques

There has been a gradual improvement in data availability. TableBuilder (Australian Bureau of Statistics 2012a) allows users to create census cross-tabulations with almost any choice of axes, rather than limiting users to tables of pre-determined content. The 11 waves of the HILDA survey have included questions on many topics. Many special-purpose longitudinal surveys are now available.

Household microsimulation techniques have also improved. Stratified sampling has been used to speed event simulation (Cumpston 2012). Immediate matching has been shown to give greater fidelity and lower runtimes than batch matching (Cumpston 2010). Parallel processing has been used in the MOSART model of 7.2 million Norwegians (Fredricksen, Knudsen & Stolen 2011). Database techniques have been used to allow unlimited numbers of kinship relationships and disease conditions (Service, Cumpston & Sarjeant 2013).

EXAMPLES OF ACTUAL USES OF MICROSIMULATION MODELS

Modelling the Australian Higher Education Contribution Scheme

The Higher Education Contribution Scheme (HECS) was introduced by the Australian government in 1989. The contribution in 1989 was $1800 for a year of study, which could become a debt to the scheme, or be paid up-front with a discount of 25%. Accumulated contributions were indexed, and became repayable at 1% of income when income reached $22,000, at 2% at $25,000 and 3% at $35,000. Voluntary repayments received a discount of 15%. Repayment rates were quickly increased, and three contribution rates, dependent on the course type, were introduced in 1998 (Jackson 2003).

The Australian Government Actuary's office constructed a microsimulation model for the scheme, simulating the buildup of HECS debt during study, and gradual repayment after earnings reached repayment thresholds. The model has been updated annually. Given the wide variations in course completion times and subsequent earnings, a microsimulation model was clearly needed.

US Congressional Budget Office CBOLT model of social security

The first version of CBOLT, completed in late 2000, was a cell-based model designed to generally mimic the methodology of the Social Security Administration's Office of the Chief Actuary (O'Hara, Sabelhaus & Simpson 2004). The second version integrated the core actuarial modules into a macro growth framework, endogenously generating wage growth and interest rates. The third version replaced most of the actuarial methods with a dynamic microsimulation model. Administrative data from the Continuous Wage History Survey gave past earnings and benefit histories for a 1% sample of all persons with social security numbers, and this was combined with other survey data to give a 0.1% sample of the US population from 1984 to 1998, including spouse links. Subsequent work added links between parents and children.

Modelling Australian compulsory superannuation contributions

Keegan (2010) used the APPSIM household microsimulation model to estimate the proportions of the population entitled to full or part age pensions. Scenarios modelled were higher labour force participation across the board, reduced disability among over-40s, higher labour force participation among primary carers of children and an increase in the mandatory superannuation contribution rate.

Modelling entry into aged care

Cumpston (2011a) derived logistic regression relationships between age care residence and age, sex, marital status and disease conditions. This allows projections of aged care requirements to be made for each local area. Developers of care facilities could use such projections to choose optimal sites and price structures for their facilities.

Modelling dwelling prices and mortgage insurance

Household microsimulation models can simulate complex interactions between individuals. For example, the behaviour of a dwelling's buyers and sellers can be modelled, with sellers reviewing their asking price on (say) a weekly basis, taking into account the offers received from buyers. Numbers of buyers and sellers can be simulated by age-based models of migration between regions, taking into account the distances between regions, the number of persons and the jobs available in the destination region. In theory, dwelling prices in each region could be endogenously generated, and estimates made of the numbers of unemployed persons with negative equity in their dwellings. The results might help determine capital requirements and underwriting limits for mortgage insurers.

CONCLUSION

As a risk measurement tool when the usual stochastic approaches are not applicable, microsimulation models provide a robust, timely and reliable method of measuring the risk involved in various scenarios being analysed. Such quality risk measurement gives a good foundation for the proper management of these types of risks.

Acknowledgements

We are very grateful for the help provided to us by Peter Martin, Paul Rodway and Browyn Loong. The suggestions made by a reviewer have improved the quality of the final paper.

Bibliography

Australian Bureau of Statistics. (2010). 1270.0.55.001 – Australian Statistical Geography Standard (ASGS): Volume 1 – Main Structure and Greater Capital City Statistical Areas, July 2011. Canberra December 23.

Australian Bureau of Statistics. (2012a). TableBuilder – user manual. http://www.abs.gov.au/websitedbs/censushome.nsf/home/tablebuilder%20manual%20contents?opendocument&navpos=240

Australian Bureau of Statistics. (2012b). 4430.0.30.002 – Microdata: Disability, Ageing and Carers, Australia, 2009. Canberra March 9.

Australian Bureau of Statistics. (2012c). 6530.0 – Household Expenditure Survey and Survey of Income and Housing, User Guide, Australia 2009–10. Canberra July 13.

Australian Bureau of Statistics. (2012d). 33010DO011_2011 Births Australia, 2011. Canberra October 25.

Australian Bureau of Statistics. (2013a). Microdata: Australian Health Survey, First Results 2011–13. Canberra.

Australian Bureau of Statistics. (2013b). Microdata: Income and Housing: Australia 2011–12. Canberra.

Australian Government Actuary. (2009). Australian Life Tables 2005-07. Downloaded from www.aga.gov.au 9/7/13.

Australian Taxation Office (ATO). (2011). Benefits, contributions, rollovers, investment earnings and final balances for each member of a self-managed superannuation fund in 07–08 and 08–09, subdivided by sex and 5-year age group, supplied in response to a Freedom of Information request, June 29.

Bellis, C., Lyon, R., Klugman, S., & Shepherd, J. (2010). Understanding actuarial mananagement: the actuarial control cycle. *Edition 2, Institute of Actuaries of Australia & Society of Actuaries*.

Caldwell, S.B., (1983). Broadening policy models: alternative strategies. Microanalytic simulation models to support social and financial policy. eds Orcutt, G.H., Merz, J., & Quinke, *H*. North-Holland, Amsterdam 1986, pp59–77.

Cumpston, J.R. (2010). Alignment and matching in multi-purpose household microsimulations. *International Journal of Microsimulation*, 3(2) 34–45.

Cumpston, J.R. (2011a). An Australian disease and long-term care microsimulation model. Paper presented to the third general conference of the International Microsimulation Association, Stockholm, June 8–10, 22 pages, available from http://www.scb.se/Grupp/Produkter_Tjanster/Kurser/_Dokument/IMA/Cumpston_paper_disease_May_10_2011.pdf

Cumpston, J.R. (2011b). New techniques for household microsimulation, and their application to Australia. PhD thesis Australian National University, available in electronic form from Chifley Library.

Cumpston, J.R. (2012). A more efficient sampling procedure, using loaded probabilities. *International Journal of Microsimulation*, 5(1) 21–30.

Fredricksen, D., Knudsen, P., & Stolen, N.M. (2011). The dynamic cross-sectional microsimulation model MOSART. Paper presented to the third general conference of the International Microsimulation Association, Stockholm, June 8–10, 8 pages.

Harding, A. (2007). Challenges and opportunities of dynamic microsimulation modelling. Plenary paper presented to the first general conference of the International Microsimulation Association, Vienna, August 21, 23 pages. International Microsimulation Association (2009) http//www.microsimulation.org

Jackson, K. (2003). The Higher Education Contribution Scheme. Downloaded 11/7/13 http://www.aph.gov.au/About_Parliament/Parliamentary_Departments/Parliamentary_Library/Publications_Archive/archive/hecs

Keegan, M. (2010). Mandatory superannuation and self-sufficiency in retirement – an application of the APPSIM dynamic microsimulation model. *Social Science Computer Review*, May 18, http://ssc.sagepub.com/content/29/1/67

Leigh, A. (2013). Battlers and billionaires. *Redback*, Collingwood, 210 pages.

Mathers, C., Vos, T., & Stevenson, C. (1999). The burden of disease and injury in Australia (1996). Cat. no. PHE 17. Canberra: AIHW. Viewed 10 July 2013 http://www.aihw.gov.au/publication-detail/?id=6442467088

O'Donoghue, C. (2001). Dynamic microsimulation: a methodological survey. *Brazilian Electronic Journal of Economics*, 4(2), 77 pages, www.microsimulationa.org/IMA/BEJE/BEJE_4_2_2.pdf downloaded 16/11/09.

Oeppen, J., & Vaupel, J.W. (2002). Broken limits to life expectancy. *Science v296* May 10 www.science.mag.org

O'Hara, J., Sabelhaus, J., & Simpson, M. (2004). Overview of the Congressional Budget Office long-term (CBOLT) policy simulation model. Congressional Budget Office, Washington DC, January.

Orcutt, G.H., Greenberger, M., Korbel, J., & Rivlin, A. (1961). Microanalysis of socioeconomic systems – a simulation study. *Harper & Brothers*, New York, xviii + 425 pages.

Percival, R. (2007). APPSIM – software selection and data structures. *NATSEM working paper no 3*, April, 15 pages.

Service, D.A., Cumpston, J.R., & Sarjeant, H.B. (2013). Synthetic baseline data for household microsimulation models for Australia and New Zealand. Abstract submitted for the International Microsimulation Association conference, Canberra, December 11–13.

Smith, D.A. (2010). Modelling. In Bellis, C., Lyon, R., Klugman, S., & Shepherd, J. (eds) *Understanding actuarial mananagement: the actuarial control cycle, edition 2*, Institute of Actuaries of Australia & Society of Actuaries, pp 227–264.

Stiglitz, J.E. (2013). *The price of inequality*. Penguin.

Summerfield, S., Freidin, S., Hahn, M., Ittak, P., Li, N., Macalalad, N., Watson, N., Wilkins, R., Wooden, M. (2012). *HILDA User Manual – Release 11*. University of Melbourne Faculty of Business and Economics, 30/11/12.

Whitelock-Jones, A. (2003). Modelling. *Understanding actuarial mananagement: the actuarial control cycle*, Bellis, C., Shepherd, J., & Lyon, R. (eds), Institute of Actuaries of Australia, pp153–172.

Wilson, T. (2009). A multistate model for projecting regional populations by Indigenous status: an application to the Northern Territory, Australia. *Environment and Planning A*, 41, 230-249.

Dr David Service FIAA, PhD, GradCert(Higher Ed)
david.service@anu.edu.au
Dr David Service is a fellow at the ANU with research interests in household microsimulation, stochastic investment modelling and economies of scale in financial services.

Dr Richard Cumpston
richard.cumpston@gmail.com
Dr Richard Cumpston is a director of Australian Projections Pty Ltd, and a school visitor at the ANU Research School of Finance, Actuarial Studies and Applied Statistics.

Hugh Sarjeant FIAA, BSc (Hons)
hugh@yirik.com.au
Hugh Sarjeant has interests in modelling financial liabilities and service requirements in aged care and superannuation, and is involved in the development of microsimulation models.

A topic of interest – how to extrapolate the yield curve in Australia

DR PETER MULQUINEY AND DR HUGH MILLER

Dr Peter Mulquiney

Dr Hugh Miller

ABSTRACT

This paper addresses key practical questions concerning the setting of interest rate assumptions in actuarial work. In particular, we look at the situation where the liabilities being valued extend beyond the term of available market instruments, and so the yield curve must be extrapolated. We find that, in Australia, the yield curve up to 2 years before the longest-dated bond can be estimated reliably. We look at international evidence and conclude there is reasonable evidence of reversion to a flat long term forward rate. However, the evidence suggests that the rate of reversion is slow with a term of about 40 years being the minimum point to reversion, and 60 years the central estimate. We also propose hedging strategies for long-dated liabilities and show how these give further insight into appropriate rates.

KEYWORDS
yield curve, long-term interest rates, uncertainty, bootstrap, linear regression, principal components analysis, hedging

AJAP 2014; 1: 13-28

1 INTRODUCTION

One of the most fundamental concepts in actuarial practice is the time value of money. For any work where future cash flows are allowed for, such as reserving or pricing, it is natural to discount to present values so that an appropriate amount of money can be set aside today to allow for future investment returns.

It is widely accepted that, for claims reserving, liabilities should be discounted using the prices of the 'risk-free' assets available in the financial markets. This means that the present value of a liability cash flow should be set equal to the market price of a basket of risk-free assets that provides a matching cash flow.

Although the principle of using risk-free rates for discounting is widely accepted, for some time there has been considerable debate on some practical aspects of the principle. This debate gained intensity following the global financial crisis of 2007–08 which saw large increases in the price of risk-free assets and correspondingly large decreases in risk-free interest rates. Issues of debate include:

- What are the best instruments to use to determine risk-free interest rates?
- Should the risk-free rate include an 'illiquidity premium'?
- What should be done when the liabilities being valued extend beyond the term of available market instruments?

In the Australian context, where we have a deep and liquid market in AAA-rated Commonwealth Government Bonds, the Australian Prudential Regulation Authority (APRA) has made it clear that it regards these bonds as the best instruments to use to determine the risk-free rate (APRA 2010; APRA GPS 320 – Discount rates). Further, for general insurance liabilities at least, APRA does not allow the inclusion of an illiquidity premium. Although GPS320 applies to general insurance work, similar requirements exist in life insurance and valuations more generally, where the risk-free rate is an input into the valuation process.

However, the issue of what should be done when the liabilities being valued extend beyond the term of available market instruments has, in our opinion, not been fully addressed in any current Australian regulations. Further, actuaries operating in the Australian market have adopted a wide range of approaches to this issue and this has led to inconsistent valuations of long-term liabilities across entities. This issue is particularly relevant for Australia, where the term of the longest dated government bond is currently around 15 years – in markets such as the US, UK and Canada government bonds are available for terms of up to 30 years or more.

Common to all approaches for the valuation of long-term liabilities is the requirement that the yield curve be extrapolated to terms beyond those available in the financial markets. The aim of this paper is to review the issues relevant to yield curve extrapolation.

Of course, the issue of yield curve extrapolation is not new. In particular it has received considerable discussion in Europe with the move towards a market-based approach to valuation under Solvency II (e.g. CEIOPS 2010; CRO Forum 2010). In presenting this work we have been heavily influenced by these discussions. In particular, we have found the papers and discussion notes by Barrie and Hibbert (Barrie & Hibbert 2008; Hibbert, 2008; Carlin 2010; Hibbert 2012) helpful and have chosen to organise this paper around the three questions they recognise as being central to the technical issue of extrapolation:

- What is the longest-maturity market forward interest rate we can estimate reliably?
- For the purposes of extrapolation, what is the ultimate very long-term 'unconditional' forward interest rate?
- What path should be set between the longest market rate and the unconditional forward rate and how many years should it take to the final level?

These issues are summarised in Figure 1.

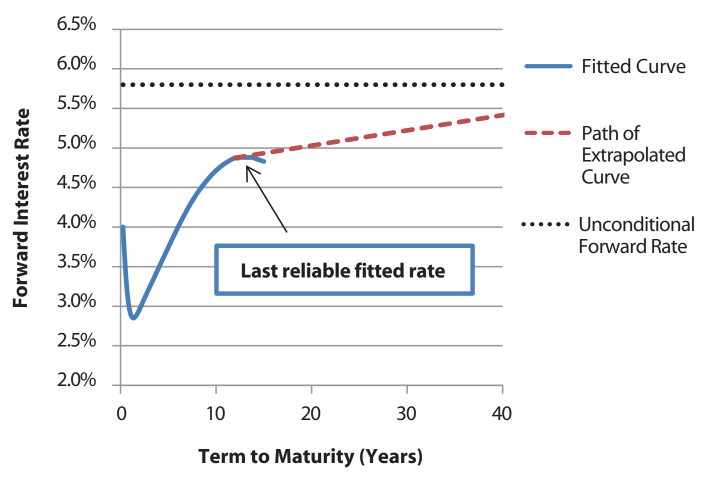

Figure 1: Yield curve extrapolation.

Our main contribution to the debate relates to the third question. In particular, we present analysis which suggests that the yield curve may only move towards its ultimate level at a relatively slow rate. The implication of this is that, for most long-tailed liabilities, insurance or otherwise, the assumed long-term 'unconditional' forward interest rate should have little impact on the valuation result. We start though, with a discussion of the first two technical questions, as they give a fuller context to the results we have found.

Before starting on the technical issues, we give a brief overview of the philosophical and regulatory issues surrounding yield curve extrapolation – these issues will influence how the technical issues should be solved.

2 PHILOSOPHICAL AND REGULATORY CONSIDERATIONS FOR YIELD CURVE EXTRAPOLATION

2.1 Philosophical considerations

Broadly speaking there are two philosophical approaches to yield curve extrapolation – one which emphasises market consistency at a point-in-time, and another which emphasises liability stability (Hibbert 2012) across time.

With the market consistency approach, the aim of yield curve extrapolation is to estimate what the market price of longer-term assets would have been on a particular day if those assets actually existed and were traded in the financial markets. The emphasis here is on producing a liability value that is equivalent to the value that would be required to transfer that liability in current market conditions.

On the other hand, those adopting an approach emphasising liability stability believe that the volatility induced by requiring market-consistent prices at a point in time is unhelpful for running a business over many years. Also they may claim that insurance liabilities are rarely traded and so the idea of identifying market value is less relevant. Some proponents of this view may be fully supportive of the idea of valuing liabilities for which hedging instruments are available with reference to the prices of those instruments. But where those instruments don't exist, or for some proponents are not sufficiently liquid, they would rather emphasise liability stability over market consistency (e.g. see CRO Forum 2010).

Each approach has its advantages and disadvantages. While an approach which emphasises liability stability may be superficially attractive, it may cause insurers to underestimate the true economic cost of writing long-term insurance contracts. It would not be in the firms', policyholders' or shareholders' future interests if long-term contracts were issued too cheaply. In addition an artificially stabilised liability valuation may also make it harder to hedge balance sheet risk over the longer term as our analysis in Section 6 seems to indicate.

On the other hand it has been suggested as part of the Solvency II debate that emphasising market price consistency may play a role in magnifying economic or financial downturns, or in other words it may add to pro-cyclicality (CRO Forum 2010).

2.2 Regulatory and professional considerations

So what approach do Australian accounting and prudential standards suggest is required? We focus on general insurance standards here, but believe that similar comments apply to life insurance and other aspects of actuarial work.

The Australian Accounting Standard AASB 1023 states:

> AASB 1023 6 – Discount Rates
> 6.1 The outstanding claims liability shall be discounted for the time value of money using risk-free discount rates that are based on current observable, objective rates that relate to the nature, structure and term of the future obligations.
> 6.1.2 Typically, government bond rates may be appropriate discount rates for the purposes of this Standard, or they may be an appropriate starting point in determining such discount rates.

So the standard is unfortunately unhelpful when observable, objective rates do not exist.

APRA's current prudential standard applicable to the valuation of insurance liabilities states:

> APRA GPS 320 – Discount rates
> 35. The rates to be used in discounting the expected future claims payments of insurance liabilities denominated in Australian currency for a class of business are derived from yields of Commonwealth Government Securities (CGS), as at the calculation date, that relate to the term of the future insurance liability cash flows for that class.
> 36. Where the term of the insurance liabilities denominated in Australian currency exceeds the maximum available term of CGS, other instruments with longer terms and current observable, objective rates are to be used as a reference point for the purpose of extrapolation. If there are no other suitable instruments or the Appointed Actuary elects to use an instrument that does not meet this requirement, the Appointed Actuary must justify the reason for using that particular instrument in the insurer's ILVR. Adjustments must be made to remove any allowances for credit risk and illiquidity that are implicit in the yields of those instruments.

The suggestion to use other instruments where the term of the liabilities exceeds the maximum available term of Commonwealth Government Bonds suggests a preference for a point-in-time market consistent valuation. Unfortunately, at present there does not appear to be a large and transparent market in alternative instruments for durations beyond 15 years. For example, the only publicly available information on bank bill swaps is for terms up to 15 years.

Finally, the Institute's PS 300, the standard applicable to general insurance valuations, mandates:

> IAA PS300 – General Insurance Business
> 8.2.2 Legislative and/or regulatory requirements may prescribe whether Claim Payments are to be discounted. The Member must consider the purpose of the valuation and document whether the future Claim Payments are to be discounted. Discount rates used must be based on the redemption yields of a Replicating Portfolio as at the valuation date, or the most recent date before the valuation date for which such rates are available.
> 8.2.3 If the projected payment profile of the future Claim Payments cannot be replicated (for example, for Classes of Business with extended runoff periods), then discount rates consistent with the intention of Paragraph 8.2.2 must be used.
> 'Replicating Portfolio' means a notional portfolio of current, observable, market-based, fixed-interest investments of highest rating, which has the same payment profile (including currency and term) as the relevant claim liability being valued.

Here the requirement to use an approach consistent with a replicating portfolio suggests a market-consistent approach should be used.

2.3 Approach taken in this paper

Although not entirely clear, it seems to us that the intention of both APRA's prudential standards and the Institute's PS300 is for yield curve extrapolation to be performed on a point in time market-consistent basis. In this paper we will approach the technical issues from this viewpoint. However, when there is uncertainty about how the market may behave at longer terms we will favour an approach that contributes to liability stability rather than an approach that may add unwarranted volatility to an insurer's balance sheet.

3 WHAT IS THE LONGEST MARKET FORWARD INTEREST RATE WE CAN ESTIMATE RELIABLY?

The question we consider in this section is: at what point along the forward interest rate yield curve should we start to extrapolate? A common starting point for extrapolation is the term of the last available market instrument. In this section we consider reasons why this may not always be appropriate. Before considering this question we will briefly discuss why we are working with the forward rate yield curve rather than one of the other possible representations of the yield curve.

3.1 Extrapolation of forward rates rather than spot rates

Yield curves can be represented in a variety of ways but those with which actuaries tend to work show either spot rates or forward rates at various terms to maturity. For this paper we have chosen to work with instantaneous (continuously compounding) forward interest rates. A reason for this is that if we were to use, say, a linear extrapolation of the spot rate curve, the change in slope would lead to a jump in the instantaneous forward rate and this would create a yield curve with arbitrage opportunities. As a general principle is seems preferable that extrapolated interest rates be arbitrage free (CRO Forum 2010).

3.2 Uncertainty in forward rate estimates

The starting point for forward rate estimation, in our case, will be the prices of Commonwealth Government Bonds. When fitting these prices to estimate the forward rates, if we were to fit the prices exactly we would end up with a very irregular (bumpy) looking forward curve. The reason for this is that that bonds of adjacent maturities can trade at prices that are slightly off-trend, as traded volumes across different maturities vary. Or, in other words, liquidity effects can lead to noise in reported bond prices. Because of this noise, a better estimate of the true underlying forward curve is obtained by imposing some smoothness on the estimated forward rates. In fact, a number of studies have shown that smooth forward rate curves perform better than bumpy ones in predicting the prices of bonds deliberately left out of the forward rate estimation process (Waggoner 1997; Bolder & Gusba 2002).

Because of the noise present in bond prices, any forward rate curve will have some estimation error associated with it. To illustrate this, we estimated 90% confidence intervals for forward rate curves from two different days over the last year (Figure 2).

Figure 2 shows that up until a term of 10 years, the 90% confidence interval is around 10–15 basis points

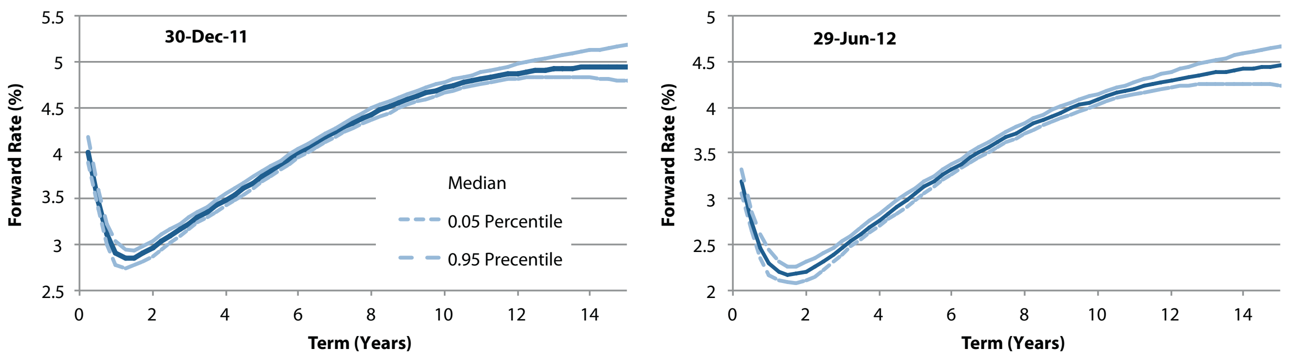

Figure 2: Estimation error in the forward rate curve.
The forward curves were estimated from the quoted prices of Australian Government Bonds on two separate days using a modified version of the Merrill Lynch Exponential Spline Method as presented in Finlay and Chambers (2008). This is the method the Reserve Bank of Australia currently uses for forward rate estimation. Error estimates were made using the bootstrap method with 1000 bootstrap samples. The forward rates that we estimated minimised the weighted least squares error in bond prices. The weights used were the reciprocal of the modified duration for each bond. The effect of this was to place less weight on longer term bonds. The specific assumption is that the variance of the bonds pricing error is proportional to its duration. Or more simply, that we expect the observed prices for these bonds to be more variable.

in width. But by year 13 it has spread to 25 basis points, and reaches about 40 basis points by year 15. Part of the reason for this increase between years 13 and 15 is the sparseness in bonds in this region. For example, the last two bonds are at terms of roughly 12 years and 15 years – a 3-year separation.

The implication of this analysis is that the uncertainty in forward rate estimates will be far greater at the long end of the curve. Indeed, similar studies from periods when the longest maturity bonds were closer to 12 years suggest that forward rates up until about 2 years prior to the longest maturity date are fairly reliable, with estimation errors increasing rapidly after that.

A further consideration that may affect the reliability of forward rate estimates at the long end is that, in most markets, the liquidity of instruments with long terms is generally lower. For example, in Australia the value on issue of 15-year Commonwealth Government Bonds is only 2% of the total value on issue of all Commonwealth Government Bonds. It may be reasonable to expect that these bonds are subject to more price noise due to their reduced liquidity.

It is worth mentioning here that the estimation error measured in Figure 2 is specific to the model we have used to estimate the forward rates. The model we used permits a great deal of flexibility in long-term forward rates. We could, however, have used an alternative model that tried to minimise estimation error at the long end, for example, by using a model that assumed the forward curve levelled out at some duration prior to 15 years. While this would reduce estimation error it would leave us no wiser as to whether the curve had actually flattened out or not, and this would, in our view, be less informative for the problem of extrapolation.

3.3 Impact on extrapolation

The analysis presented in this section suggests that we should not be overly reliant on forward rate estimates made at the long end of the fitted forward curve, in particular the last 2 years of the observable range. At present it appears that terms of around 13 years would be an appropriate point to start an extrapolation in Australia.

Although we are recommending not placing too much reliance on the forward rates estimated at the long end, if in the process of extrapolation there was a significant deviation between the fitted curve and the extrapolated curve some care must be taken. For example, it would seem to us to be important to check that the prices of the longest-dated bonds were still estimated to within an acceptable tolerance level when the extrapolated curve was used.

With these considerations in mind some other questions are posed by the results of Figure 2. In particular, given the curves are still rising at a term of 13 years:

- At what level should the forward rates ultimately be extrapolated to?
- How quickly should they get there?

The first of these questions is answered in the next section.

4 FOR THE PURPOSES OF EXTRAPOLATION, WHAT IS THE ULTIMATE VERY LONG-TERM 'UNCONDITIONAL' FORWARD INTEREST RATE?

The rational expectations hypothesis of the term structure of interest rates suggests, in its more general form, that long-term forward rates are the sum of the expected future short-term interest rate plus a constant 'term premium' that varies only by term and not over time. Taking this to be true, the process of determining the ultimate very long-term 'unconditional' forward interest rate (UFR) would involve making estimates of what these two components are. Unfortunately, things are not so simple.

The rational expectations hypothesis has been shown to be inconsistent with the empirical evidence. In particular, a number of studies have shown that long-term interest rates are significantly more volatile than would be expected if the rational expectations hypothesis held. In addition, studies have also shown that long-term rates respond to information that would reasonably be thought to influence short-term rates only. These inconsistencies result from time-varying term premia (Shiller 1990).

It is now commonly accepted that time-varying term premia account for most of the variability in long-term forward rates (Kim & Orphanides 2005; Finlay & Chambers 2008). A fundamental question to be answered then is how far along the term structure do these term premia variations extend (Hibbert 2008)?

In the following section we discuss this question and other important considerations needed to project expected future short-term interest rates and term premia out to very long terms.

4.1 The difference between term and time

Before we begin this section it is important to understand the distinction between term and time. On any given day there is a yield curve of spot rates or forward rates that varies by term to maturity, or equivalently, duration. The yield curve itself will also evolve over calendar time as the prices of bonds of different durations change over time. The problem of determining the UFR involves determining what the forward rate is at very long durations on a specified day. In this section we will discuss the idea that the UFR can be determined as the sum of the expected future short-term interest rate plus the term premium applicable at very long durations on the specified day. Here the expected future short-term interest rate is the markets expectation on the specified day of what short-term interest rates will be at a maturity date long into the future.

4.2 Expected short-term interest rate

When projecting expected future short-term interest rates to maturity dates long into the future it is reasonable to expect that our ultimate rate should only change in response to fundamental changes in the structure of the economy and should not be affected by short-term economic changes.

With this in mind, it is usual (e.g. Barrie & Hibbert 2008; CRO Forum 2010) to separate the problem of setting the expected short-term interest rate into one of determining:

- expected future inflation;
- expected future real short-term interest rate.

In relation to expected future inflation, the Reserve Bank of Australia has been very successful in targeting inflation and entrenching low and stable inflation expectations for at least the last 15 years (Finlay & Chambers 2008). It seems reasonable then to adopt the mid-range of the bank's current CPI target of 2–3% as our future inflation expectation. The issue of forecasting inflation rates was dealt with in more detail in a paper by Miller (2010).

For expected real short-term interest rates a typical approach has been to look at historical averages for real cash returns across several countries (e.g. CEIOPS 2010; Barrie & Hibbert 2008). The underlying assumption is that, in the long-run, real interest rates should not differ substantially across economies which are at a similar stage of economic development. As part of Solvency II, the QIS 5 assumed the expected real rate of return to be 2.2% (CEIOPS 2010). This figure was determined by reference to a study by Dimson et al. (2000) looking at bond returns over the second half of the 20th century for 12 major economies, including Australia. A similar study by Barrie and Hibbert (2008) found a median estimate of 1.8% from 16 developed countries over the period 1970–2007. In light of these studies a figure of around 2% seems reasonable to us.

4.3 Term premia

The definition of term premia that is most useful in the current context is that they are the difference between the forward rate and the expectation of the future short-term interest rate. Term premia have a number of causes:

- Investors demand a term premium for locking into long-term investments. In this case the term premium acts as compensation for holding long term bonds whose value will fluctuate in the face of interest rate uncertainty, exposing the holder to mark to market losses. Term premia in this case are often referred to as risk premia and will be positive.
- Alternatively, demand for long-term

government securities from large institutional investors such as insurance companies and pension funds can drive down long-term forward rates because these long-term bonds offer a closer match to liabilities and are less risky investments to these investors. Such forces can cause term premia to be negative. This phenomenon is often called 'term preference'.
- Additionally, convexity effects also cause term premiums to decrease at longer maturities. Fixed income securities have positive convexity. This means that the capital gains and losses from equal sized interest rate swings are not matched – the gains will be greater than the losses. This effect can theoretically cause very-long duration bonds to trade at higher prices (lower yields).

Term premia tend to increase quickly up to around 2 years duration and reach a relatively flat level by about 10 years. Term premia at around 10 years are on average of the order of 1% to 2% and at this duration appear to be dominated by risk premia effects. In many currencies, starting around 15 to 20 years, there is a decline in the forward rates which may be of the order of 1% by year 30. This decline is primarily due to the term preference of pension and life insurance companies, but also due the increased convexity of long-maturity bonds.

However, this is a stylised view of term premia and ignores their important time-varying behaviour. This behaviour is discussed in the following sub-section.

4.3.1 The time-varying behaviour of term premia

Model-based techniques exist which allow the decomposition of forward rates into expected future short-term interest rates plus term premia. These models may incorporate survey data about analysts' forecasts of short-term interest rates. Those who use these models acknowledge that they are complex and difficult to calibrate and that the results should not be interpreted too precisely. However they do appear to provide important insight into the time-varying nature of term premia.

The following figure shows the results from one such analysis (Kim & Wright 2005) on US Treasuries. The figure shows how forward rates at different terms can be broken down into the underlying expected short rate and term premium components. The results show the daily variation.

Some relevant observations from Figure 3 include:
- Term premia are particularly volatile. There can be large month-to-month swings and over the course of a year it is not uncommon for the term premia to change by 1% - 2%. The changes in term premia at year 5 are often mirrored in the changes at year 10.
- Term premia tend to increase when yields are rising and tend to decrease when yields are falling. These changes sometimes appear as over-reactions to tightening or loosening monetary policy.
- There are some noticeable trends in term premia. In particular term premia declined over the early 1990s until around mid-1998. This decline corresponded to a general decline in US inflation expectations over this period.

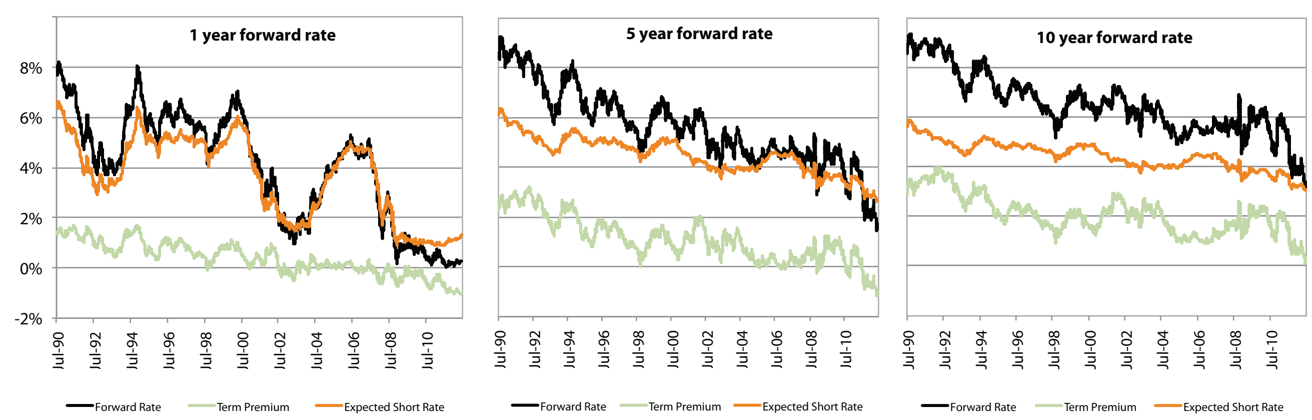

Figure 3: Term premia estimated for US Treasuries.
The data in this figure was obtained from the US Federal Reserve website and was estimated using the method outlined in the paper by Kim and Wright (2005).

- Towards the end of 2011 term premia have declined dramatically at all durations. For some durations the term premia are negative suggesting investors are willing to pay a premium for government bonds. This is a result of the 'flight to quality' in response to the bad news from Europe.

For our purposes, the most important observation from Figure 3 relates to the variability of the term premia. We see that while most of the variability of the 1-year forward rate is explained by variability in the expected short rate, by year 10, most of the variability is explained by variability in the term premia. This is, in part, because short-term interest rates are expected to mean revert over time. Or in other words, expectations of future short-term interest rates are relatively fixed by durations of 10 years or more. On the other hand, the term premia appear to display little (or no? – it is not certain from the study) mean reversion with term.

Similar studies have been performed on Australian Commonwealth Government bonds with similar results. Indeed there is a very close correlation between Australian and US term premia suggesting that they are driven by global rather than country specific factors (Finlay & Chambers 2008).

4.3.2 Implications for UFR determination

As discussed earlier, the term structure for a particular date depends only on the prices of bonds quoted on that date. These prices in turn depend on the economy wide supply and demand for risk free credit on that particular day. And as we have seen in Figure 3 supply and demand changes may result in large swings in term premia from day to day and month to month.

An unresolved question is how far down the term structure do these day to day and month to month swings extend? Unfortunately there are no empirical results which can help us with this question. It certainly seems possible that the volatility in term premia could extend out across the entire term structure if bonds of very long maturities were actually traded and particularly, if this volatility was driven by changes in the risk premia required for investing long term. But this is mere speculation and in the words of Hibbert (2008), 'how far is it reasonable to adjust the 50-year, 100-year and 1000-year forward rate?'

So the pragmatic response is to set the long-term term premium assumption to some assumed stable level. This may not be correct, but it steers away from introducing potentially unwarranted volatility onto an insurer's balance sheet.

4.4 Assumptions for components of the UFR

Pulling together the arguments of this section, at the present time a UFR of about 5.8% seems reasonable for Australia. The components of this estimate are detailed in the following table.

Table 1: Components of the unconditional forward rate for Australia in 2012.

Component	Rate
Expected future short-term interest rate	
Expected future inflation	2.5%
Expected future real interest rate	2.0%
	4.5%
Term premium	
Risk Premium	1.5%
Convexity adjustment	–0.2%
	1.3%
Unconditional forward rate	**5.8%**

As discussed earlier the component of the UFR that is most uncertain is the term premium. A value of 1.3% was chosen using the following considerations:

- Using the data from Figure 3, we calculate that the average term premium on 10 year bonds in the US, since the stabilisation of US inflation expectations from around 2007, is 1.7%.
- Work by the Reserve Bank of Australia has shown that since 1997, term premia on 5-year forward rates have been around 0.5% to 1% lower than those in the US (Finlay & Chambers 2008). However in that work, estimated short-term rates that were about 1.5% higher than those assumed here.
- While there are **term preference** effects in markets such as the US and UK beyond terms, of around 15 years, we are unsure if it is appropriate to allow for them in the UFR for Australia. It has been suggested that, at very long terms, there are no liabilities to hedge anyway so the hedging activities of institutional investors are unlikely to bid down term premia beyond say 100 years (Carlin 2010).
- For the convexity adjustment we have relied on the work of the CRO Forum (2010) which showed that the convexity premium decreased about 0.2% from year 10 to year 30 (where it reached a minimum). Other estimates have been in the range of 0.4% (Barrie & Hibbert 2008).

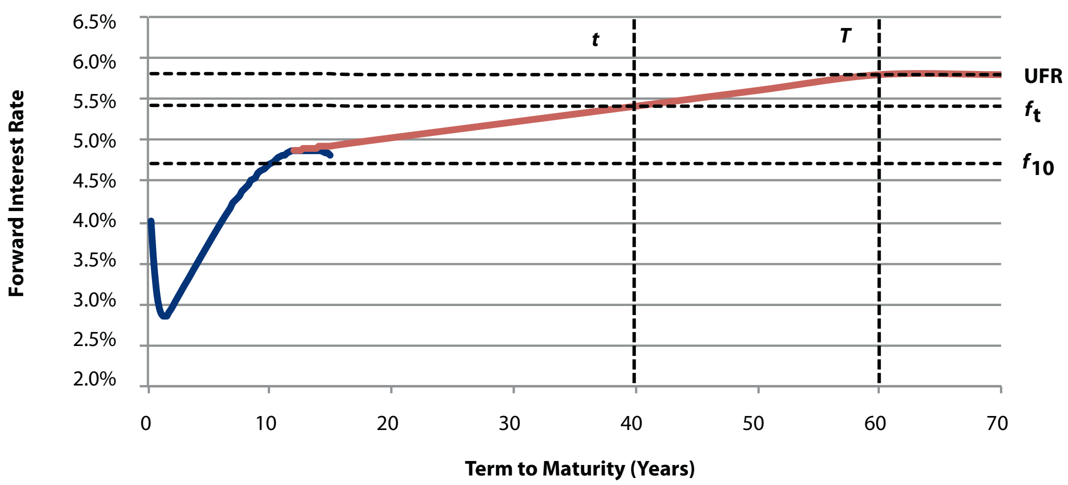

Figure 4: Linear extrapolation model.

- Given uncertainties in both the risk premium component and the convexity adjustment, a combined figure of around 1.3% seems reasonable.

As a reasonableness check on our adopted UFR of 5.8%, we note that the average 10-year forward rate in Australia has been 5.8% since 1998. Strictly, speaking, to compare our UFR to the recent average 10 year forward rate we need to remove the convexity adjustment – since by our definition this adjustment is zero at year 10 – giving a rate of 6.0%, which is slightly above the recent 10 year forward rate average.

It should be apparent from the discussion in this section that there is some uncertainty about what an appropriate UFR assumption should be – a figure anywhere in the range 5.4% –6.2% could be reasonably justified at this point in time. For most insurance liabilities including long-tailed classes, this uncertainty will only have a material impact if the UFR is reached relatively quickly in an extrapolation. This issue of the speed of reversion to the UFR is discussed in the next section.

5 WHAT PATH SHOULD BE SET BETWEEN THE LONGEST MARKET RATE AND THE UNCONDITIONAL FORWARD RATE?

5.1 Introduction

We now move on to the final, and arguably most important, aspect of yield curve extrapolation: what path should be set between the longest duration market rate and the unconditional forward rate and how long should we take to reach it?

Note there is an implicit assumption in the previous question – that the extrapolated rate should approach to the UFR. There is a possibility that the extrapolated curve should be flat, or equivalently that it takes a very long time to approach the UFR. The international evidence offered below suggests that this is not the case.

The speed of return to the UFR is of practical importance; in the Australian context, if it returns quickly (by duration 20 years say), then long-tailed liabilities will be considerably more stable over time than if it returns slowly (e.g. by 100 years).

Throughout this section we will focus primarily on a linear extrapolation between observed bond rate and the UFR, illustrated schematically in Figure 4. This is partly out of a desire for simplicity, but is also consistent with much of current practice in Australia and New Zealand (including the approach mandated by the New Zealand Treasury for government accounting work). The conclusion of Section 5 will give some consideration to non-linear patterns of extrapolation.

This pattern of linear extrapolation to the UFR should be applicable to any shape of yield curve. If the longest term reliable forward rate is above the UFR, then the extrapolation slope is downward. A forward rate below the UFR would result in an upward slope of extrapolation.

5.2 Insights from countries with longer dated government bonds

5.2.1 Background and data sources

In Australia our longest bond term is around 15 years. In countries such as the US, UK and Canada however government bonds are available at terms of 30 years or longer. So it seems to us forward rate curves from these countries would be a good place to start to understand how the Australian yield curve should be extrapolated, at least in the range 15–30 years.

The central banks from the US, UK and Canada make public their historical estimates of the forward rate curve for government bonds in their respective countries. We have used these forward rate curves for our analysis. We have relied only on the forward rate curves available from 1 January 1998. We believe this is a good starting point for the current era of monetary policy characterised by low and stable inflation expectations.

5.2.2 A regression based analysis

Consider the simple linear regression equation:

$$f_{s+t} = \alpha_{s:t} + \beta_{s:t} f_s \quad (1)$$

Here the intercept $\alpha_{s:t}$ allows for the UFR as well as any (fixed) term premiums across the yield curve, and $\beta_{s:t}$ is the linear dependence of the forward rate at term $s+t$ on the forward rate at term *s*. If the schematic in Figure 4 is correct, then $\beta_{s:t}$ should be a good indicator of progress towards the UFR; if the slope is close to 1, this implies the forward rate at $s+t$ is still moving in sync with the rate at s, so no reversion to the UFR has taken place. Note too that if the slope was consistently close to 1 as t grows, this would be strong evidence against any reversion to the UFR. Conversely, if the slope is close to zero then the forward rate at $s+t$ is largely independent of the rate at s, suggesting that it has reverted to a constant level. Values in between can indicate the rate of progress, so a value of 0.5 indicates a place on the extrapolation halfway along the reversion.

To illustrate, we regress f_{20} against f_{10} using the UK data. The resulting fit is $f_{20} = 0.78\% + 0.74 f_{10}$.

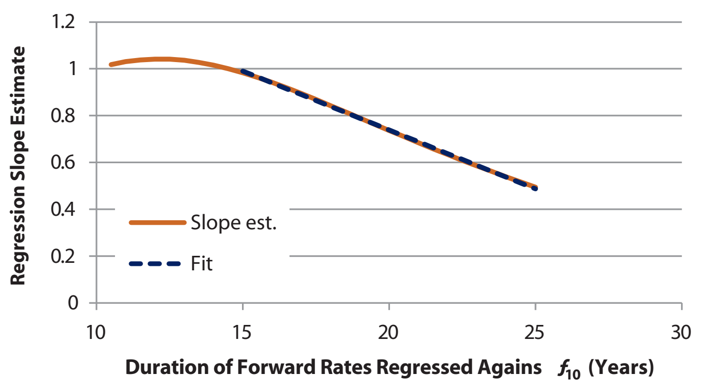

Figure 5: Slope (β) coefficients for UK forward rates regressed against f_{10}. So the slope estimate of 0.74 for f_{20} (discussed above) can be seen at duration 20 on the plot.

The slope parameter suggests a heavy relationship between the two, but some evidence of mean reversion. We can then hold s =10 fixed, and vary t between 0 and 20 to build up a more complete picture. The results are presented in Figure 5.

The results here are surprisingly clear. The coefficients are relatively stable until term 15, where they begin to decrease at a rate very close to linear. Extrapolating that rate of decrease would lead to the slope equalling zero (and thus reaching the UFR) at duration 35 years. We discuss the uncertainty around this estimate below.

Similar analyses can be performed on the USA and the Canadian forward rate data. A summary of these results, together with a rough measure of uncertainty is shown in Table 2. The uncertainty on the point estimates is determined by bootstrapping the historical data series and refitting the linear trends.

Table 2: Regression results for the linear extrapolation model (Equation 1). Standard errors in brackets.

Country	Duration decay starts	Duration when reach UFR	95% confidence interval
US	10	82	(55, 168)
UK	15	34	(31, 40)
Canada	10	41	(35, 47)

The large degree of uncertainty, particularly in the slow decay seen in the USA data, is reflected in the relatively large confidence intervals. However, the results from this analysis appear clear enough to allow us to conclude that:

- there is reasonable international evidence for reversion to a UFR
- the reversion starting at somewhere between duration 10 and 15 is plausible
- the reversion is quite slow – it is not complete before duration 30, and in some cases is considerably longer.

5.2.3 Principal components analysis of yield curves

Principal components analysis (PCA) is a means of determining what yield curve shapes account for the majority of the variation in the yield curve over time. PCA has been used in a variety of finance contexts, including yield curves (Litterman & Scheinkman 1991), foreign exchange (Avellaneda & Zhu 1997) and equities (Gourrieroux et al 1997; Laloux et al 1999). The most important feature of a PCA is that the shape of the components shows the typical movement of the yield curve around the mean. This is illustrated in Figure 6.

The two curves shown in the chart are the expected shapes under two different hypotheses. The flat line is the component expected when most movement in the yield curve is attributable to parallel shift – the short and long durations move in equal amounts. The reversion shape is what would be expected when the forward rates return to a UFR type average in the long term (here after 25 years); the shape of the curve means that there is more movement in the short part of the curve and little at the high durations, where it remains closer to the UFR. By examining the shape of principal components in other countries we can gain significant insight as to the speed for which forward rates might return to the UFR.

Data for the analysis was the same as that used in the previous regression analysis.

Figure 7 shows the Canadian principal component analysis. The first heavy dark component dominates the results – it accounts for over 80% of the total variation. This curve is relatively flat, which we can interpret to represent shifts in the yield curve that are close to parallel. In addition to the primary principal component, there are other shapes that explain other types of movements in the yield curve. The second curve (8% of variation) is a flex shape for when medium durations move relative to short and long rates, and the third curve (6%) is a flex primarily at the short durations. The other 6% of total yield curve variation is accounted for by the remaining principal components, which are not shown here.

The shape and heavy importance of the first component (10 times as much variation explained compared with any other component) suggests that the **bulk of the yield curve movement is explained by parallel shift**. This fact is consistent with the observations of Section 4.3.1 about the behaviour of time-varying term premia. That said, there is some evidence that the component is decaying from about duration 16 onwards; this can be seen as the slight decrease observable in the blue curve from this term. While this decay is slight, it is statistically significant. We examine these trends in further detail at the end of this section.

Similar decompisitions can also be carried out for the UK and US. In all cases we observe that each of the leading components show some decay towards zero at high durations. Although slight, these trends are statistically significant (if we take bootstrap samples and refit, the trends persist), and the (assumed linear) slope of the downtrend can be used to estimate the point at which the curves reach zero, which is analogous to reaching the UFR. The zero point corresponds to reaching the UFR because this is the point at which there is no (or relatively little) movement in the forward rates. This provides a further check to the results in

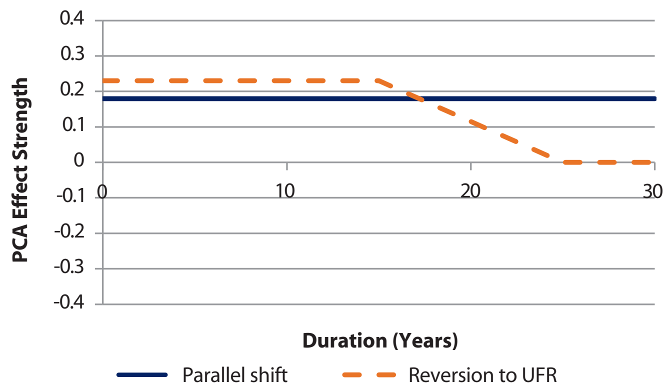

Figure 6: Expected principal component curves under reversion to UFR and parallel shift hypotheses.

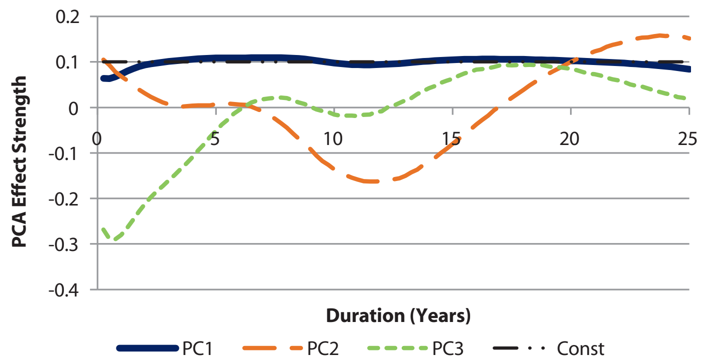

Figure 7: Principal components analysis of Canadian yield curve.

Table 3: Results for extrapolation of decay in leading principal component.

Country	Starting duration	Slope	95% CI	Duration when reach UFR	95% CI
Canada	16	–0.0023	(–0.0032,–0.0015)	64	(46,101)
UK	15	–0.0026	(–0.0037, –0.0018)	84	(60,122)
USA	17	–0.0021	(–0.0027, –0.0016)	110	(87,147)

Section 5.2.2. The results are summarised in Table 3 (see page 23) – we have given estimates of the slope, the duration at which the curves reach zero using this slope, as well as 95% confidence intervals on these estimates (using bootstrap resampling).

The results between countries are strikingly similar, despite the natural differences expected between countries and the sensitivity of the duration estimates to small changes in slope. In particular, all three confidence intervals include the duration range (87,100).

While there are issues with the analysis (such as the statistical validity of performing regression on a PCA), the results do point to a flat shape to the forward rate curve, with very slow decay to the UFR.

5.2.4 Reconciling the regression and PCA results and other methodological considerations

The PCA approach gives a reversion period significantly longer than the regression analysis. While we do not completely understand the reasons for this, we favour the results of the regression analysis as is the most direct test of the problem of extrapolation: given the last reliable estimate of the forward rate what projected future forward rate will give me the least prediction error? Further the PCA approach only analyses the behaviour of the first principal component and ignores other sources of variation. However, we like the insights provided by PCA – that the bulk of the yield curve movement is explained by parallel shift and we are pleased that it provides results broadly consistent with the regression analysis.

5.3 Discussion of results and other research

The concept of extrapolation to a long term rate over a long period is not unique to this paper. Two relevant examples are:

- QIS 5 undertaken for Solvency II suggested that convergence to an UFR should be reached between terms of 70–90 years; while
- the method presented by Barrie and Hibbert (2008) suggested convergence should occur by terms of around 100 years.

The approach used by Barrie and Hibbert to arrive at their 100-year convergence period is not in the public domain, but as far as we can tell it seems to involve making subjective judgements about what forward rate volatility should be at longer terms based on what it was in measured forward rates up to about terms of about 15 years. An approach obviously requiring a significant amount of judgement, but from what we have seen, it doesn't appear entirely unreasonable and it has also been endorsed by the CRO Forum (2010).

Both the QIS 5 and the Barrie and Hibbert approach involve extrapolations that are not linear – the paths chosen have a decaying rate – which would lead to a somewhat longer time to reach the UFR.

The latest indications are that the Solvency II guidelines will require a significant earlier convergence term than was initially suggested in QIS 5 – possibly somewhere between 10 and 40 years after the last reliable (in this case liquid) term – but this seems to be driven in part by some of the stability concerns outlined in Section 2.

The choice of non-linear paths is partly a tension between simplicity and aesthetics, given that identifying the true shape involves considerable uncertainty. We believe the linear path is reasonable, but recognise that there are more sophisticated techniques out there such as:

- the Smith-Wilson technique (described in CEIOPS 2010)
- the use of the Nelson-Siegal functional form (described in Barrie & Hibbert 2008).

Discussions of the pros and cons of each of these methods can be found in the cited references. An advantage of a method such as Nelson-Siegal is that it allows one to match the slope of the fitted curve at the start of the extrapolation. This would be an advantage if the slope of the fitted curve was rising relatively rapidly or if it was turning downwards and away from the UFR at its final point.

5.4 Final thoughts on the extrapolation path

The analysis of yield curves in countries with longer-dated securities indicates that it may revert to the UFR somewhere between duration 40 and 100 years, with about term 60 years a reasonable recommendation from our results. This is mostly based on the average of the regression results, but gives some small weight to the longer PCA estimates. While there may be instances where quicker convergence to the ultimate forward rate could be justified – for example where the fitted forward curve was continuing to rise rapidly near its final terms – the practice of simply assuming quick convergence, say within a 5 or 10 year period, cannot be justified.

6 IMPLICATIONS FOR HEDGING BALANCE SHEET RISK

It is possible to use a duration-matching strategy to hedge risks that are beyond the longest term assets available, if you are allowed to take short positions in assets, and have an estimate for the risk-free rate beyond the longest term asset. We have not seen this hedging concept in the actuarial literature (in which case it is overdue), but we would expect it to be common knowledge amongst traders that hedge long positions.

Given that this hedging approach requires an extrapolated yield curve, a natural question is whether hedging performance is better or worse if the extrapolated yield curve is assumed to revert quickly to the UFR?

We have attempted to answer this question by testing a duration matching hedging strategy under two alternative assumptions:

1. that the forward rates reach ultimate level very slowly, say by a term of 80 years (the point-in-time market consistent approach); or
2. that the forward rates converge to ultimate very quickly; say by a term of 20 years (the stability approach).

The tests have been performed using historical data from the Australian Government Bond market.

6.1 Details of the hedging strategy

The hedging strategy is as follows. Suppose we have a distant fixed liability L_T, payable at time T in the future. Suppose further that there are two shorter dated risk-free zero coupon bonds available in the market with durations s and t, with $s < t < T$. Then it is possible to match the present value and the (modified) duration of the liability by shorting A_s and going long A_t. In formulae, it is possible to choose loadings a and b, with $a < 0 < b$, such that:

(2) $$aPV(A_s) + bPV(A_t) = PV(L_T)$$

and:

(3) $$Dur(aA_s + bA_t) = Dur(L_T)$$

Here Dur is the modified duration measuring the price sensitivity to changes in interest rate y:

(4) $$Dur(L_T) = -\frac{1}{PV(L_T)}\frac{\partial PV(L_T)}{\partial y}$$

Thus we have duration matched our long liability using shorter term assets. This portfolio can be rebalanced at regular intervals to ensure duration remains matched. We make the following comments:

- There is no reason why only two assets have to used, or that the assets are zero coupon bonds – these are simplifying assumptions for illustration and testing purposes. In fact, more complex duration-matched portfolios are likely to exhibit superior properties (e.g. better convexity properties).
- The concept of duration matching for hedging is a type of immunisation (Redington 1952). In standard immunisation the portfolio is chosen to ensure the convexity of the assets is greater than that of the liabilities, giving small (second order) profits in the presence of parallel shifts in the yield curve. In our situation the convexity of the assets is always less, exposing possible second order loss. This may be a secondary issue in practice, if:
  - there is enough non-parallel shift to swamp these second order effects; and
  - the existence of term premiums (usually higher yields for longer bonds) means that this strategy collects some of this long term premium, partly offsetting convexity losses.
- This approach results in large opposing long and short positions in government bonds. This obviously comes with a number of associated costs which we have not evaluated and so we are unsure how attractive this hedging strategy would actually be in practice.
- The hidden assumption in the equations above is that $PV(L_T)$ is known. Of course, this requires extrapolation of the risk-free rates between terms t and T. Thus the effectiveness of the hedging strategy depends (somewhat) on the appropriateness of the extrapolation assumptions. We discuss this further below.

6.2 Australian market examples

We have run a number of historical experiments testing this strategy. The general setup is:

- The distant liability of $100 falls due in 20 years.
- We invest $PV(L_T)$ in the market to cover this liability.
- We invest in only 4 and 10 year zero coupon bonds, returning the risk free market rate.
- We rebalance every quarter. This involves closing the positions in the 3.75 and 9.75 year bonds and opening a new position in 4 and 10 year bonds, still matching the duration of the liability (19.75 after the first quarter etc).
- After any quarter we can calculate the asset/liability ratio to assess how well we have hedged balance sheet movements. The variability of the asset/liability ratio around 100% indicates how good the hedging is.
- We close the position after 10 years of managing (as at that time we can perfectly match the liability by investing in a 10 year ZCB, crystallising any gain/loss).

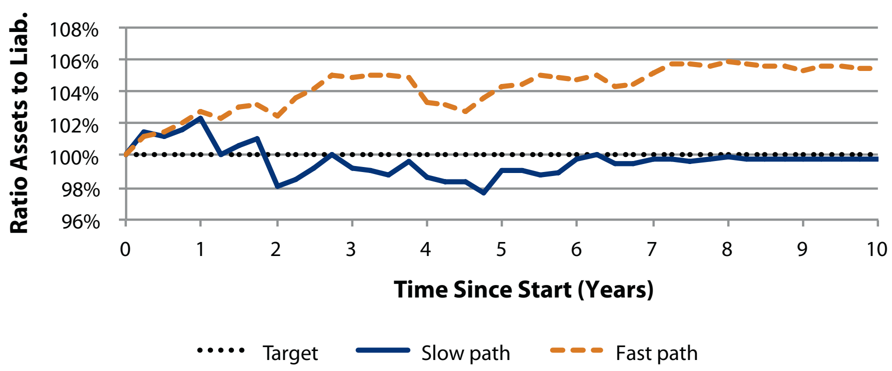

Figure 8: Comparison of long liability hedging performance, starting June 1995.

Note that the assumption of quick reversion to the UFR has the effect of reducing the modified duration of the liability, because a portion of its value is stabilised by the path assumption. A linear reversion to the UFR between terms 10 and 20 gives a modified duration of the liability of 14.9 years. Thus this assumption requires a less extreme short-long position of bonds, compared with the slow reversion.

The dark solid line in Figure 8 shows the inflated value of our assets if we use a flat forward rate (a reasonable proxy for any slow path of return) beyond 10 years, starting at time 30 June 1995. That is, we set the forwards between term 10 and 20 at the same level as the 10-year forward rate. The performance of this strategy from 1995 is remarkably good, despite large changes in interest rates over the date range. At the end of the management period a tiny loss is made (about 20c on the $100), but over the course of the 10 years of management the implied value hugs the target 100% value fairly tightly. One way to quantify this is the average absolute departure from 100% (that is, the average distance between blue and orange lines), which in this case is 0.8%.

We can then compare the hedging performance if we use an alternative yield curve beyond ten years. Here we test a mean reversion assumption that the forward rate returns to 6.0% (consistent with that recommended in Section 4.4 without allowance for the convexity assumption) linearly over the duration years 10 through 20. The resulting hedging performance is the dashed line in the figure.

The performance is markedly worse, with an average absolute difference of 4.2%. In this case the eventual result is a substantial profit (due to the falling rates in the late 90s), but this is not the aim of a hedging program. The slow path seems far superior at reducing interest rate risk. We note however, that while we have assumed a constant UFR over the 10 year period, in practice one could imagine a higher UFR being adopted in the first few years from 1995, with some lower rate adopted towards the end of the 1990s when lower inflation expectations had been cemented. However, we tried a number of alternative scenarios anticipating such a UFR revision without changing our general conclusion – that the market consistent approach gave a better result.

We can run the experiment over different historical periods. In Table 4 we summarise results for June 1998, 2000, and 2002 starting points. This Table shows the average distance between the each hedging performance line in the figures and the target 100% line. The slow path outperforms in all except the year 2000 according to this metric.

Table 4: Comparison of hedging performance using average absolute difference metric.

Year (30-Jun)	Slow path	Fast path	Ratio
1995	0.8%	4.2%	5.64
1998	0.6%	0.9%	1.44
2000	0.9%	0.7%	0.73
2002	0.7%	1.1%	1.57

To summarise, the differences in the later experiments are less stark than the 1995 result, with the slow path clearly better in 1998 and 2002, but more ambiguous in 2000. The extreme difference in the first date range is not completely unexpected; the period 1995 to 1998 was unique due to the very strong decrease in yields from 9% to about 6%; this is the type of situation where a good hedging strategy is most critical.

We believe this approach to hedging appears fairly legitimate from both a theoretical and historical data perspective. It gives hedging error of less than 1% in all experiments. The fast path is inferior, giving further evidence that a slower reversion to UFR is closer to 'truth'. Investing only in the 10-year bond (a common strategy in practice – 'hold the longest term bond possible') by contrast would result in significant losses in the first and last experiments, where a slide in the 10 year forward rate is observed.

7 CONCLUSIONS AND RECOMMENDATIONS

We can summarise our findings relatively succinctly:

- **The yield curve up to 2 years before the longest-dated bond can be estimated reliably.** For the last year or so, the noise and method of fit can cause significant (relative) error.
- **There is reasonable international market evidence for reversion to a flat long term forward rate.** This rate is reached via extrapolation from the end of the observable yield curve.
- **The rate of reversion is slow.** We believe term 40 is about the minimum point to reversion based on the bond markets examined, with a central estimate closer to term 60. This conclusion rests on the assumption that the unconditional forward rate has been stable over the period 1998 to 2012.
- **Linear path reversion is plausible, with other approaches possible.** The regression analysis and PCA results generally support a linear decay to the UFR; however much of the uncertainty regarding reversion period applies equally to reversion shape. Non-linear paths may have implications for term at which the long term rate can be reached.
- **Long-term risk-free hedging is possible, at least for moderate term extrapolations of the yield curve.** The lack of long-dated risk-free assets does not mean no hedging strategies exist for long dated liabilities. The methods proposed here appear plausible and give reasonable results on historical data for cash flows up to 20 years in the future. However, they do require a reasonable estimate of yields beyond the observable range.

We believe that these results make significant contributions to actuarial assumption setting.

Finally, we make two comments. First, the current Australian accounting and prudential standards, along with the Actuaries Institute's PS300, are not entirely clear, in our opinion, in stating the philosophy that should be used for yield curve extrapolation. We understand that there are some current efforts being made by the Institute to address yield curve extrapolation.

Second, the conclusion that the rate of reversion is slow means there is an inevitable tension between market consistency and liability stability; while it may be attractive from a stability standpoint to have a yield curve that quickly reverts to an assumed long-term average, we believe the valuation would be incorrect if appropriate long term risk free assets existed. The tension deserves further consideration in the actuarial community.

Bibliography

Avellaneda, M., & Zhu, Y. (1997). An E-arch model for the term structure of implied volatility of FX options, *Int. J. Theor. Appl. Finance.*

Barrie & Hibbert (2008). A framework for estimating and extrapolating the term structure of interest rates. Exposure Draft.

Bolder, D.J., & Gusba, S. (2002). Exponentials, Polynomials, and Fourier Series: More yield curve modelling at the Bank of Canada. Bank of Canada working paper.

Bolder, D.J., Johnson, G., & Metzler, A. (2004). An empirical analysis of the Canadian Term structure of zero-coupon interest rates. Bank of Canada working paper.

Carlin, S. (2010). Thoughts on QIS5 yield curves. Barrie+Hibbert Insights.

CEIOPS (2010). QIS5 Risk free interest rates. Extrapolation Methods.

CRO Forum (2010). CRO Forum Best Practice Paper – Extrapolation of Market Data.

Dimson, E., Marsh, P., & Staunton, M. (2000). Risk and return in the 20th and 21st. *Business Strategy Review*, 11, (2), 1-18.

Finlay, R., & Chambers, M. (2008). A term structure decomposition of the Australian yield curve. Reserve Bank of Australia Research Discussion Paper.

Gourrieroux, Ch., Scaillet, O., & Szafarz, A., (1997). *Econometrie de la Finance*, Ed. Economica, Paris.

Gurkaynak, R.S., Sack, B., & Wright, J.K. (2007). *The U.S. Treasury Yield Curve: 1961 to the Present.* Finance and economic discussion series, Federal Reserve Board, Washington D.C.

Hibbert, J. (2008). Market-consistent valuation of ultra long-term cash flows, Barrie+Hibbert Insights.

Hibbert, J. (2012). Alternative views on extrapolated yield curves: A fundamental question remains unanswered. Barrie+Hibbert Insights.

Kim, D.H., & Orphanides, A. (2005). Term structure estimation with survey data on interest rate forecasts. Federal Reserve Board Finance and Economics Discussion Series.

Kim, D.H., & Wright, J.H. (2005). An arbitrage-free three-factor term structure model and the recent behaviour of long-term yields and distant-horizon forward rates. Federal Reserve Board Finance and Economics Discussion Series.

Laloux, L., Cizeau, C., Bouchaud, J.P., & Potters, M. (1999). Noise dressing of financial correlation matrices, *Physical Review Letters*, 83**,** 1467.

Litterman, R., & Scheinkman J. (1991). Common factors affecting bond returns, *J. Fixed Income* 55-61.

Miller, H. (2010). Towards a better inflation forecast. Paper for the 2010 General Insurance Seminar, Institute of Actuaries Australia.

Redington, F.M. (1952). Review of the Principles of Life Office Valuations. *Journal of the Institute of Actuaries*, 78, 286-340.

Shiller, R.J. (1990). The term structure of interest rates in *Handbook of Monetary Economics Volume 1*. B.M. Friedman & F.H. Hahn (eds).

Svensson, Lars E.O. (1994). Estimating and Interpreting Forward Rates: Sweden 1992-4, National Bureau of Economic Research Working Paper #4871.

Waggoner, D. (1997). Spline methods for extracting interest rate curves from coupon bond prices, Federal Reserve Bank of Atlanta, Working Paper series, 97–10.

Dr Peter Mulquiney
peter.mulquiney@taylorfry.com.au
Dr Peter Mulquiney is an actuary at Taylor Fry with over 10 years experience providing advice to general insurers and government.

Dr Hugh Miller
hugh.miller@taylorfry.com.au
Dr Hugh Miller is an actuary at Taylor Fry. His research interests include actuarial assumption setting, large dataset analysis and predictive analytics .

Hedging long-dated interest rate derivatives for Australian pension funds and life insurers

DR KEVIN FERGUSSON AND PROF ECKHARD PLATEN

Dr Kevin Fergusson

Prof Eckhard Platen

ABSTRACT

Many pension funds and life insurers seek to hedge their exposure to low interest rates using long-dated interest rate derivatives. This paper extends an approach of Platen and Heath 2006 to price and hedge long-dated interest rate derivatives using a combination of Australian cash, bonds and equities and under a variety of market models. The results show the models under which the lowest cost hedge is achieved.

KEYWORDS

growth optimal portfolio, benchmark approach, long-dated zero coupon bonds, minimal market model

1. INTRODUCTION

In this article we describe the hedging strategy used to replicate the payoffs of interest rate derivatives at maturity, in particular payoffs corresponding to a zero coupon bond and a swaption. Crucially, we calculate the cost of purchasing a hedge portfolio at the outset of hedging the zero coupon bond (ZCB), which ensures that the ZCB payoff is affordable by the hedge strategy with high probability. We compute the costs of hedging ZCBs across the considered market models and all ZCB terms to maturity and identify the best performing models.

Pricing and hedging of long-dated derivative payoffs remains a difficult and unsolved problem in financial and actuarial industries. Previous work on hedging long-dated zero coupon bonds was published in Platen 2006 and Bruti-Liberati and Platen 2010 where they make use of the growth optimal portfolio (GOP) which is maximising the logarithmic utility of expected terminal wealth. Platen employs the minimal market model (MMM) stock index dynamics to obtain low-cost replicating hedge portfolios for zero coupon bonds. The analysis in the current article extends this strategy to market models with stochastic interest rates. It is our intention to demonstrate that for market models employing the MMM discounted GOP the costs of hedging are significantly cheaper than for those employing the Black-Scholes form of the discounted GOP.

The market models examined here are specified by the stochastic differential equation (SDE) of the short rate r_t and the SDE of the discounted GOP $\bar{S}_t^{\delta_*}$. The short rate models considered are the deterministic short rate model, where the short rate is known for all times:

(1) $$r_t = r(t)$$

the Vasicek short rate model described by Vasicek 1977:

(2) $$dr_t = \kappa(\bar{r} - r_t)dt + \sigma dZ_t$$

the Cox-Ingersoll-Ross (CIR) short rate model described by Cox et al. 1985:

(3) $$dr_t = \kappa(\bar{r} - r_t)dt + \sigma\sqrt{r_t}dZ_t$$

and the 3=2 short rate model described by Ahn and Gao 1999:

(4) $$dr_t = (pr_t + qr_t^2)dt + \sigma r_t^{3/2}dZ_t$$

The discounted GOP models which are considered are the Black-Scholes model, equivalently the lognormal stock price model, employed by Black and Scholes 1973:

(5) $$d\bar{S}_t^{\delta_*} = \bar{S}_t^{\delta_*}\psi^2 dt + \bar{S}_t^{\delta_*}\psi dW_t$$

and the minimal market model described by Platen 2001:

(6) $$d\bar{S}_t^{\delta_*} = \bar{\alpha}_t dt + \sqrt{\bar{\alpha}_t \bar{S}_t^{\delta_*}}dW_t$$

where $\bar{\alpha}_t = \bar{\alpha}_0 \exp(\eta t)$. Here Z_t and W_t are independent Wiener processes, *r(t)* is the realised value of the short rate at time *t* and $\bar{r},\ \kappa,\ \sigma,\ p,\ q,\ \psi,\ \bar{\alpha}_0$ and η are constants.

The cash account B_t is the accumulated value at time *t* of $1 deposited at initial time zero and we have the formula:

(7) $$B_t = \exp\left(\int_0^t dr_s\right)$$

The GOP $S_t^{\delta_*}$ is obtained by multiplying the cash account B_t by the discounted GOP $\bar{S}_t^{\delta_*}$. The growth rate of the GOP is equal to the drift term of the SDE of the logarithm of the GOP, which for models involving the Black-Scholes discounted GOP is $g_t = r_t + \frac{1}{2}\psi^2$ and for models involving the MMM discounted GOP is $g_t = r_t + \frac{1}{2}\bar{\alpha}_t/\bar{S}_t^{\delta_*}$.

For a given contingent claim H_T with maturity $T \in (0, \infty)$, it has been shown in Platen and Heath 2006 that the minimal possible price $V_t^{\delta_{H_T}}$ for a replicating hedge portfolio satisfies the real world pricing formula:

$$V_t^{\delta_{H_T}} = E_t\left(\frac{S_t^{\delta_*}}{S_T^{\delta_*}} H_T\right) \tag{8}$$

where E_t denotes the real world conditional expectation given the real world probability measure. Furthermore, the GOP S^{δ_*} is taken here as the numeraire or benchmark. The numeraire is the portfolio having maximal growth rate and is approximated well by diversified equity market indices such as the S&P/ASX 200 Total Return Index or S&P 500 Total Return Index.

Under our considered market models the real world pricing formula (8) gives the price of a ZCB as:

$$P(t,T) = E_t\left(\frac{B_t}{B_T}\right) E_t\left(\frac{\bar{S}_t^{\delta_*}}{\bar{S}_T^{\delta_*}}\right) \tag{9}$$

For the deterministic model of the short rate we have:

$$E_t\left(\frac{B_t}{B_T}\right) = \exp\left(-\int_t^T r(s)ds\right) \tag{10}$$

For the Vasicek model of the short rate we have from Vasicek 1977:

$$E_t\left(\frac{B_t}{B_T}\right) = A(t,T)\exp(-r_t B(t,T)) \tag{11}$$

where:

$$B(t,T) = \frac{1 - \exp(-\kappa(T-t))}{\kappa} \tag{12}$$

and:

$$A(t,T) = \exp\left((\bar{r} - \frac{\sigma^2}{2\kappa^2})(B(t,T) - T + t) - \frac{\sigma^2}{4\kappa}B(t,T)^2\right) \tag{13}$$

For the CIR model of the short rate we have from Cox et al. 1985:

$$E_t\left(\frac{B_t}{B_T}\right) = A(t,T)\exp(-r_t B(t,T)) \tag{14}$$

where:

$$A(t,T) = \left(\frac{h\exp(\frac{1}{2}\kappa(T-t))}{\kappa\sinh\frac{1}{2}h(T-t) + h\cosh\frac{1}{2}h(T-t)}\right)^{2\kappa\bar{r}/\sigma^2} \tag{15}$$

$$B(t,T) = \frac{2\sinh\frac{1}{2}h(T-t)}{\kappa\sinh\frac{1}{2}h(T-t) + h\cosh\frac{1}{2}h(T-t)} \tag{16}$$

and:

$$h = \sqrt{\kappa^2 + 2\sigma^2} \tag{17}$$

For the 3/2 model of the short rate we have from Ahn and Gao 1999:

$$E_t\left(\frac{B_t}{B_T}\right) = \frac{\Gamma(\gamma_1 - \alpha_1)}{\Gamma(\gamma_1)}\left(\frac{2}{\sigma^2 y(t,r_t)}\right)^{\alpha_1} M(\alpha_1, \gamma_1, \frac{-2}{\sigma^2 y(t,r_t)}) \tag{18}$$

where:

$$y(t, r_t) = \frac{r_t}{p}(\exp((T-t)p) - 1) \tag{19}$$

$$\alpha_1 = -\left(\frac{1}{2} - \frac{q}{\sigma^2}\right) + \sqrt{\left(\frac{1}{2} - \frac{q}{\sigma^2}\right)^2 + \frac{2}{\sigma^2}}$$

$$\gamma_1 = 2\left(\alpha_1 + 1 - \frac{q}{\sigma^2}\right)$$

Here M is the confluent hypergeometric function given by:

$$M(\alpha, \gamma, z) = \sum_{n=0}^{\infty} \frac{(\alpha)_n}{(\gamma)_n} \frac{z^n}{n!} \tag{20}$$

and $\Gamma(x) = \int_0^\infty u^{x-1} \exp(-u)$ is the gamma function. For the Black-Scholes discounted GOP we have:

$$E_t\left(\frac{\bar{S}_t^{\delta_*}}{\bar{S}_T^{\delta_*}}\right) = 1 \tag{21}$$

whereas for the MMM discounted GOP we have from Platen and Heath 2006:

$$E_t\left(\frac{\bar{S}_t^{\delta_*}}{\bar{S}_T^{\delta_*}}\right) = 1 - \exp\left(-\frac{1}{2}\bar{S}_t^{\delta_*}/(\varphi_T - \varphi_t)\right) \tag{22}$$

where $\varphi_t = \frac{1}{4}\bar{\alpha}_0(\exp(\eta t) - 1)/\eta$. Thus various combinations of (11), (14), (18), (21) and (22) inserted into (9) give explicit formulae for the real world prices of ZCBs under each considered market model.

2. DESCRIPTION OF METHODOLOGY

In respect of a derivative security a hedging strategy is a trading strategy involving a portfolio of hedge securities whose value at a prescribed payoff date is intended to replicate the value of the derivative security.

When the market values of securities are driven by a deterministic short rate and stochastic discounted GOP then we have only one random factor in our market and we can hedge a suitable derivative security using a managed self-financing portfolio of cash (the savings account) and the GOP. The value of the hedge portfolio can be written as:

$$V_t^{(\pi)} = \delta_t^{(0)} B_t + \delta_t^{(1)} S_t^{\delta_*} \tag{23}$$

where $\delta_t^{(0)}$ is the number of units of the cash account and $\delta_t^{(1)}$ is the number of units of the GOP account at time $t \in [0, T]$. The respective fractions invested at time $t \geq 0$ are $\pi_t = (\pi_t^{(0)}, \pi_t^{(1)})$ with $\pi_t^{(0)} = \delta_t^{(0)} B_t / V_t^{(\pi)}$ and $\pi_t^{(1)} = 1 - \pi_t^{(0)} = \delta_t^{(1)} S_t^{\delta_*} / V_t^{(\pi)}$. We have some flexibility in our choice of hedge securities and we could have used instead the savings account and futures on the GOP, for example.

When the market values of securities are driven by a stochastic short rate and a stochastic discounted GOP then we have two random factors in our market and we can hedge any derivative security using a managed portfolio of cash B_t, the GOP index $S_t^{\delta_*}$ and, for instance, a $(T-t)$-year zero coupon bond $F(t, T)$. The value of the hedge portfolio π can be written as:

$$V_t^{(\pi)} = \delta_t^{(0)} B_t + \delta_t^{(1)} S_t^{\delta_*} + \delta_t^{(2)} F(t, T) \tag{24}$$

where $\delta_t^{(0)}$ and $\delta_t^{(1)}$ describe numbers of units as before, and $\delta_t^{(2)}$ is the number of units of the T-maturity zero coupon bond at time $t \in [0, T]$.

The cost C_t at time t of hedging a derivative since initial time 0 is equal to the cost of the derivative at time t less any gains from trading the hedge portfolio. We write:

$$C_t = V_t^{\delta_{H_T}} - \int_0^t \delta_u^{(0)} dB_u - \int_0^t \delta_u^{(1)} dS_u^{\delta_*} = V_t^{\delta_{H_T}} - \int_0^t dV_u^{(\pi)} \tag{25}$$

where $V_t^{\delta_{H_T}}$ is the value of the derivative at time *t* and $V_t^{(\pi)}$ is the value of the hedge portfolio at time *t*.

This equation can be rewritten as:

$$C_t = V_t^{\delta_{H_T}} - (V_t^{(\pi)} - V_0^{(\pi)}) \tag{26}$$

$$= V_0^{(\pi)} + (V_t^{\delta_{H_T}} - V_t^{(\pi)}) \tag{27}$$

and we can see that the cost of hedging can be expressed alternatively as the cost of the hedge portfolio at outset, namely $V_0^{(\pi)}$, plus additional funds needed at time *t* to purchase the derivative in excess of the value of the hedge portfolio.

At the payoff date *T* the cost of hedging is:

$$C_T = V_T^{\delta_{H_T}} - \int_0^T dV_u^{(\pi)} \tag{28}$$

Because we are interested in the real world price of hedging, as given in (8), we consider the benchmarked cost of hedging, computed as:

$$\hat{C}_T = \frac{C_T}{S_T^{\delta_*}} = \hat{V}_T^{\delta_{H_T}} - \int_t^T d\hat{V}_u^{(\pi)} = \hat{V}_t^{(\pi)} + \hat{V}_T^{\delta_{H_T}} - \hat{V}_T^{(\pi)} \tag{29}$$

According to (8) the average of the benchmarked costs of hedging performed over a large number of backtests ought to approximate the real world price of the derivative with payoff H_T.

Given a fully specified model with known parameters, we backtest hedging of the derivative over the time interval [0; *T*] by setting the $n-1$ rebalancing times:

$$t_1 < t_2 < \dots < t_{n-1}$$

satisfying $0 = t_0 < t_1$ and $t_{n-1} < t_n = T$. The hedge portfolio $V^{(\pi)}$ is adjusted at the rebalancing times and is computed iteratively using the formula:

$$V_{t_i}^{(\pi)} = \delta_{t_{i-1}}^{(0)} B_{t_i} + \delta_{t_{i-1}}^{(1)} S_{t_i}^{\delta_*} + \delta_{t_{i-1}}^{(2)} F(t_i, T) \tag{30}$$

for $i = 1, 2, \dots, n$ with initial condition:

$$V_0^{(\pi)} = V_0^{\delta_{H_T}} \tag{31}$$

where, for $i = 1, 2, \dots n - 1$, the numbers of units held in the GOP and the ZCB at time t_i are computed as:

$$\begin{aligned} \delta_{t_i}^{(1)} &= \frac{\partial}{\partial S_s^{\delta_*}} V_s^{\delta_{H_T}}(r_s, S_s^{\delta_*})\Big|_{s=t_i} - \delta_{t_i}^{(2)} \frac{\partial}{\partial S_s^{\delta_*}} F(s,T)\Big|_{s=t_i} \\ \delta_{t_i}^{(2)} &= \frac{\partial}{\partial r_s} V_s^{\delta_{H_T}}(r_s, S_s^{\delta_*})\Big|_{s=t_i} \Big/ \frac{\partial}{\partial r_s} F(s,T)\Big|_{s=t_i} \end{aligned} \tag{32}$$

and the number of units held in the cash account at time t_i is computed as:

$$\delta_{t_i}^{(0)} = \left(V_{t_i}^{(\pi)} - \delta_{t_i}^{(1)} S_{t_i}^{\delta_*} - \delta_{t_i}^{(2)} F(t_i, T)\right) / B_{t_i} \tag{33}$$

3. ASSESSING THE PERFORMANCE OF A HEDGING STRATEGY

A perfect hedge strategy is one for which:

$$C_t = V_0^{(\pi)} \tag{34}$$

for all times $t \in [0, T]$. That is to say, the hedge portfolio replicates the value of the derivative over the life of the hedging strategy.

However, perfect hedging is not possible for many reasons and we are interested in strategies which generate the payoff at expiry date T, with 'minimum' cost.

Therefore, for a given market model, a given data set and a given ZCB term to maturity we compute the benchmarked costs of hedging a ZCB at maturity over all possible periods within the data set. From this the pth percentile of the set of benchmarked costs is computed. The best hedge strategy is derived from the market model which gives the minimum percentile benchmarked cost of hedging. Consequently, our task in this article is to compare the percentile benchmarked costs of hedging across all mentioned market models.

4. HEDGE SECURITIES

As stated earlier, when the market values of securities are driven by a stochastic short rate and a stochastic discounted GOP, we have two random factors in our market and we can theoretically hedge a suitable derivative security using a managed portfolio of cash, the GOP index and a 10-year coupon bond, for example. Because liquidity is essential for any hedge strategy we would choose to hedge using a managed portfolio of cash, S&P/ASX 200 Index Futures and 10Y Government Bond Futures, the futures contracts being tradeable on the Sydney Futures Exchange. For a US-based strategy then we would choose to hedge using a managed portfolio of cash, S&P 500 Index Futures and 10Y US Treasury Bonds which are more liquid than the corresponding Australian securities.

Table 1: Maximum likelihood estimates of model parameters fitted to Australian data Jan 1980 – Dec 2012.

Model	Parameters	Standard errors	Log likelihood
Vasicek	$\bar{r}$ = 7.3274%	3.1031%	1356.1581
	κ = 15.936%	10.5493%	
	σ = 2.7237%	0.097636%	
CIR	$\bar{r}$ = 7.1221%	2.514%	1460.1961
	κ = 13.7747%	9.735%	
	σ = 7.5906%	0.27192%	
3/2	p = 0.02%	7.558%	1563.0387
	q = 20.38%	98.9268%	
	σ = 0.7679	0.02746	
Black-Scholes	ψ = 18.5831%	3.5553%	450.9212
MMM	$\bar{\alpha}_0$ = 0.050444	0.0062622	447.5751
	η = 0.036237%	0.7008%	

Table 2: Maximum likelihood estimates of model parameters fitted to US data 1871–2012.

Model	Parameters	Standard errors	Log likelihood
Vasicek	$\bar{r}$ = 4.2294%	0.80023%	399.7019
	κ = 16.2953%	5.3703%	
	σ = 1.5384%	0.099592%	
CIR	$\bar{r}$ = 4.1078%	1.1421%	427.8116
	κ = 9.2540%	3.8668%	
	σ = 6.4670%	0.40761%	
3/2	p = 3.8506%	4.2499%	406.2713
	q = 87.7908%	119.0438%	
	σ = 2.0681	0.13425	
Black-Scholes	ψ = 17.7297%	5.9087%	-267.4135
MMM	$\bar{\alpha}_0$ = 0.010028	0.0023389	-264.6433
	η = 4.6841%	0.28769%	

5. MARKET DATA AND FITTING THE MODELS

The Australian data set consists of monthly series of 90-day bank bill rates, 10Y government bond yields and the All Ordinaries Accumulation index from January 1980 to December 2012. Each short rate model and discounted GOP model has an explicit formula for the transition density function and this has allowed us to *t* the models to the historical data using maximum likelihood estimation (MLE).

The maximum likelihood estimates (MLEs) of the parameters of all models fitted to Australian data are shown in Table 1.

Because of the relative shortness of the Australian data series we have chosen to consider a longer data set for the sake of obtaining results for long-dated contracts.

The data set used for backtesting has been the annual series of US 1Y deposit rates, 10Y treasury bond yields and S&P Composite Stock Index from 1871 to 2012, shown in Chapter 26 of Shiller [1989] and subsequently updated on Shiller's website http://aida.wss.yale.edu/ shiller/data/chapt26.xls. The 141-year length of this data series makes it a most useful series for analysing the hedging of long-dated ZCBs because we are able to backtest any given hedge strategy over the large term to maturity of the ZCB. Also, because there are 10Y bond yields accompanying the 1Y deposit rates and stock index values we are able to construct and backtest a hedge portfolio which immunises against movements in both the stock index and short rate. The MLEs of the parameters of all models fitted to US data are shown in Table 2.

The backtests of the hedging strategies were performed using in-sample estimation of parameters. Of course in reality one would backtest a hedge strategy using out-of-sample parameter estimates but by employing in-sample estimates any poorly performing model is readily falsified.

6. COSTS OF HEDGING ZCBs UNDER DETERMINISTIC SHORT RATE MODELS

We present the costs of hedging ZCBs under deterministic short rate models.

In Table 3 the percentile benchmarked costs of hedging ZCBs of various terms to maturity and percentiles are shown for the deterministic short rate and Black- Scholes discounted GOP model.

In Table 4 the percentile benchmarked costs of hedging ZCBs of various terms to maturity and

Table 3: Percentile costs of hedging ZCBs under a deterministic short rate and Black-Scholes discounted GOP based on US data 1871–2012.

Term to maturity of ZCB	99th percentile	95th percentile	90th percentile	85th percentile	80th percentile
1Y	0.99473	0.99256	0.99	0.98445	0.97523
2Y	0.98889	0.98546	0.97681	0.96485	0.95167
3Y	0.983	0.97716	0.96162	0.94172	0.92744
4Y	0.97626	0.96988	0.94909	0.92018	0.89415
5Y	0.96928	0.96266	0.94193	0.89956	0.86139
7Y	0.95547	0.94773	0.91388	0.86023	0.81143
10Y	0.93309	0.92335	0.86542	0.80451	0.75071
15Y	0.88791	0.85433	0.79846	0.71259	0.65539
20Y	0.80645	0.78216	0.7149	0.65548	0.58733
25Y	0.71779	0.67893	0.64009	0.58695	0.49921
30Y	0.60825	0.59331	0.5553	0.50713	0.43517
40Y	0.40573	0.39554	0.3682	0.34926	0.32832
50Y	0.24457	0.2396	0.23292	0.23088	0.21696

Table 4: Percentile costs of hedging ZCBs under a deterministic short rate and MMM discounted GOP based on US data 1871–2012.

Term to maturity of ZCB	99th percentile	95th percentile	90th percentile	85th percentile	80th percentile
1Y	0.99473	0.99256	0.99	0.98445	0.97523
2Y	0.98889	0.98546	0.97681	0.96485	0.95167
3Y	0.98302	0.97716	0.96162	0.94172	0.92744
4Y	0.97649	0.96952	0.94906	0.92013	0.89414
5Y	0.96964	0.95494	0.92906	0.89953	0.86139
7Y	0.9499	0.91168	0.88088	0.84826	0.81098
10Y	0.89031	0.83149	0.79108	0.75737	0.73169
15Y	0.74213	0.67106	0.61237	0.57877	0.56427
20Y	0.64319	0.49652	0.4638	0.41004	0.38913
25Y	0.52576	0.3897	0.31444	0.29971	0.29182
30Y	0.38523	0.29224	0.22849	0.22021	0.21614
40Y	0.16308	0.12845	0.11989	0.11917	0.11351
50Y	0.068659	0.065715	0.05967	0.057967	0.054139

Table 5: Percentile costs of hedging ZCBs under a Vasicek short rate and Black-Scholes discounted GOP based on US data 1871–2012.

Term to maturity of ZCB	99th percentile	95th percentile	90th percentile	85th percentile	80th percentile
1Y	1.0013	0.98871	0.98544	0.98271	0.97847
2Y	0.99264	0.9739	0.96921	0.9637	0.9467
3Y	0.97404	0.95791	0.94726	0.93674	0.9253
4Y	0.96205	0.93717	0.92472	0.91721	0.90395
5Y	0.9355	0.91767	0.90768	0.89717	0.87567
7Y	0.88889	0.87213	0.86141	0.83933	0.82671
10Y	0.8127	0.79847	0.78845	0.77038	0.74663
15Y	0.67968	0.67295	0.66559	0.649	0.63895
20Y	0.56756	0.55927	0.55355	0.54078	0.52354
25Y	0.46926	0.46426	0.4608	0.44949	0.43604
30Y	0.38976	0.38519	0.38057	0.37245	0.36232
40Y	0.26676	0.2641	0.26108	0.25494	0.24617
50Y	0.17771	0.17666	0.17453	0.17411	0.16753

Table 6. Percentile costs of hedging ZCBs under a Vasicek short rate and MMM discounted GOP based on US data 1871–2012.

Term to maturity of ZCB	99th percentile	95th percentile	90th percentile	85th percentile	80th percentile
1Y	1.0051	0.99447	0.9874	0.98391	0.97682
2Y	1.0014	0.98466	0.97472	0.96521	0.94847
3Y	0.99603	0.96746	0.95574	0.93713	0.92878
4Y	0.98577	0.95668	0.93736	0.92271	0.9063
5Y	0.95988	0.94028	0.92605	0.91038	0.88276
7Y	0.91948	0.89163	0.88089	0.85599	0.83274
10Y	0.87379	0.81513	0.77244	0.75994	0.7546
15Y	0.6831	0.65907	0.64572	0.63447	0.6218
20Y	0.54225	0.529	0.51811	0.50842	0.49915
25Y	0.42601	0.41738	0.40556	0.39134	0.38235
30Y	0.3573	0.33452	0.3185	0.29529	0.2832
40Y	0.22491	0.18573	0.17552	0.16199	0.14944
50Y	0.12039	0.1056	0.090466	0.084736	0.076245

Table 7: Percentile costs of hedging ZCBs under a CIR short rate and Black-Scholes discounted GOP based on US data 1871–2012.

Term to maturity of ZCB	99th percentile	95th percentile	90th percentile	85th percentile	80th percentile
1Y	1.0001	0.98956	0.98631	0.98302	0.97828
2Y	0.99272	0.97622	0.97054	0.9648	0.94575
3Y	0.97723	0.96202	0.95097	0.93704	0.92666
4Y	0.96774	0.94217	0.93011	0.92037	0.90335
5Y	0.9425	0.92712	0.91128	0.90274	0.87842
7Y	0.90316	0.88562	0.86989	0.84718	0.82948
10Y	0.83723	0.82484	0.80931	0.78952	0.75835
15Y	0.72898	0.72066	0.70421	0.67551	0.65364
20Y	0.62868	0.61946	0.60705	0.58503	0.56702
25Y	0.53487	0.52786	0.5194	0.50455	0.47967
30Y	0.45474	0.44968	0.44021	0.42824	0.40364
40Y	0.32668	0.32162	0.31693	0.30465	0.29126
50Y	0.23041	0.22724	0.22549	0.22036	0.20748

Table 8: Percentile costs of hedging ZCBs under a CIR short rate and MMM discounted GOP based on US data 1871–2012.

Term to maturity of ZCB	99th percentile	95th percentile	90th percentile	85th percentile	80th percentile
1Y	1.0033	0.99453	0.98864	0.98375	0.97739
2Y	1.0006	0.98628	0.97576	0.96508	0.94935
3Y	0.99508	0.971	0.95914	0.93814	0.92994
4Y	0.9886	0.96117	0.94117	0.92378	0.90995
5Y	0.96787	0.94869	0.93089	0.90962	0.88337
7Y	0.93207	0.90247	0.89243	0.86901	0.83723
10Y	0.8765	0.82884	0.79637	0.7777	0.75824
15Y	0.68994	0.67036	0.65193	0.6415	0.62889
20Y	0.55194	0.53916	0.52856	0.5136	0.50601
25Y	0.43752	0.42795	0.41449	0.40905	0.39516
30Y	0.3682	0.34461	0.33042	0.30899	0.29755
40Y	0.23746	0.19375	0.17902	0.17116	0.16419
50Y	0.13086	0.11298	0.099064	0.093498	0.083863

Table 9: Percentile costs of hedging ZCBs under a 3=2 short rate and Black-Scholes discounted GOP based on US data 1871–2012.

Term to maturity of ZCB	99th percentile	95th percentile	90th percentile	85th percentile	80th percentile
1Y	0.99498	0.99158	0.98821	0.98359	0.97967
2Y	0.98935	0.98128	0.97437	0.96548	0.94832
3Y	0.98466	0.96941	0.96053	0.94512	0.92963
4Y	0.97618	0.95459	0.94238	0.93253	0.91251
5Y	0.96757	0.94845	0.93463	0.91505	0.88228
7Y	0.93604	0.91496	0.89838	0.86453	0.83134
10Y	0.89952	0.88432	0.84687	0.81012	0.78084
15Y	0.85513	0.83077	0.77158	0.71565	0.67678
20Y	0.81023	0.77288	0.71775	0.6346	0.58291
25Y	0.76106	0.71113	0.6639	0.57766	0.51403
30Y	0.70388	0.65474	0.60165	0.52649	0.45337
40Y	0.59249	0.54902	0.50104	0.44657	0.37547
50Y	0.51234	0.47976	0.44259	0.39518	0.33001

Table 10: Percentile costs of hedging ZCBs under a 3/2 short rate and MMM discounted GOP based on US data 1871–2012.

Term to maturity of ZCB	99th percentile	95th percentile	90th percentile	85th percentile	80th percentile
1Y	0.99785	0.99392	0.98956	0.98369	0.97805
2Y	0.995	0.98708	0.97712	0.96662	0.94873
3Y	0.99639	0.97634	0.96354	0.94541	0.93078
4Y	0.98765	0.96792	0.95093	0.93281	0.91262
5Y	0.9722	0.95564	0.94524	0.92132	0.89381
7Y	0.95036	0.92299	0.90941	0.86955	0.83559
10Y	0.88323	0.86718	0.82947	0.80948	0.77803
15Y	0.7512	0.71025	0.68236	0.67631	0.65572
20Y	0.59172	0.55556	0.54667	0.53749	0.52085
25Y	0.47171	0.44381	0.43052	0.41659	0.4072
30Y	0.37711	0.35252	0.34447	0.32434	0.3118
40Y	0.24959	0.20988	0.20229	0.19247	0.1888
50Y	0.14498	0.12691	0.12073	0.11236	0.10697

percentiles are shown for the deterministic short rate and MMM discounted GOP model.

For hedging ZCBs with terms to maturity shorter than 10 years, the BS discounted GOP model and the MMM discounted GOP model perform similarly. At or beyond a 10-year term to maturity the MMM discounted GOP model significantly outperforms the BS discounted GOP model. For example, hedging a 50Y ZCB at the 99th percentile incurs a cost of 0.068859 under the MMM discounted GOP model, which is about a quarter of the corresponding cost of 0.24457 under the BS discounted GOP model.

7. COSTS OF HEDGING ZCBs UNDER VASICEK SHORT RATE MODELS

We present the costs of hedging ZCBs under Vasicek short rate models.

In Table 5 the percentile benchmarked costs of hedging ZCBs of various terms to maturity and percentiles are shown for the Vasicek short rate and Black-Scholes discounted GOP model.

In Table 6 the percentile benchmarked costs of hedging ZCBs of various terms to maturity and percentiles are shown for the Vasicek short rate and MMM discounted GOP model.

For hedging ZCBs with terms to maturity at or less than 15 years, the BS discounted GOP model and the MMM discounted GOP model perform similarly. However, beyond a 15-year term to maturity the MMM discounted GOP model significantly outperforms the BS discounted GOP model. In particular, the cost of hedging a 50Y ZCB at the 99th percentile is 0.12039 under the MMM discounted GOP model, which is about two-thirds of the corresponding cost of 0.17771 under the BS discounted GOP model.

8. COSTS OF HEDGING ZCBs UNDER CIR SHORT RATE MODELS

We present the costs of hedging ZCBs under CIR short rate models.

In Table 7 the percentile benchmarked costs of hedging ZCBs of various terms to maturity and percentiles are shown for the CIR short rate and Black-Scholes discounted GOP model.

In Table 8 the percentile benchmarked costs of hedging ZCBs of various terms to maturity and percentiles are shown for the CIR short rate and MMM discounted GOP model.

For terms to maturity shorter than 15 years, the BS discounted GOP model and the MMM discounted GOP model provide similar costs of hedging ZCBs. However, at or beyond a 15-year term to maturity, the MMM discounted GOP model significantly outperforms the BS discounted GOP model. For example, the cost of hedging a ZCB is significantly reduced for a 50 year term to maturity at the 99th percentile, the cost being 0.13086 under the MMM discounted GOP model which is about a half of the corresponding cost of 0.23041 under the BS discounted GOP model.

9. COSTS OF HEDGING ZCBs UNDER 3/2 SHORT RATE MODELS

We present the costs of hedging ZCBs under 3/2 short rate models.

In Table 9 the percentile benchmarked costs of hedging ZCBs of various terms to maturity and percentiles are shown for the 3/2 short rate and Black-Scholes discounted GOP model. In Table 10 the percentile benchmarked costs of hedging ZCBs of various terms to maturity and percentiles are shown for the 3/2 short rate and MMM discounted GOP model.

For ZCB terms to maturity at or shorter than 10 years the BS discounted GOP model and the MMM discounted GOP model provide similar costs of hedging ZCBs. But beyond a term to maturity of 10 years the MMM discounted GOP model significantly outperforms the BS discounted GOP model. For example, the cost of hedging a 50Y ZCB at the 99th percentile is 0.14498 under the MMM discounted GOP model which is roughly a quarter of the corresponding cost of 0.51234 under the BS discounted GOP model.

10. CONCLUSIONS ON THE ZCB HEDGING PERFORMANCES

In Figure 1 the 99th percentile costs of hedging ZCBs of varying terms to maturity are graphed. Each model for which the discounted GOP is modelled by the MMM has significantly cheaper costs of hedging long-dated ZCBs, that is, ZCBs with maturities beyond 15 years. In particular, we find that among the considered market models having a stochastic short rate, the Vasicek short rate and MMM discounted GOP model provides the cheapest hedging strategy for long-dated ZCBs. In Sections 11 to 14 we compare model performances on hedging swaptions.

Figure 1: Percentile costs of hedging ZCBs of varying terms to maturity.

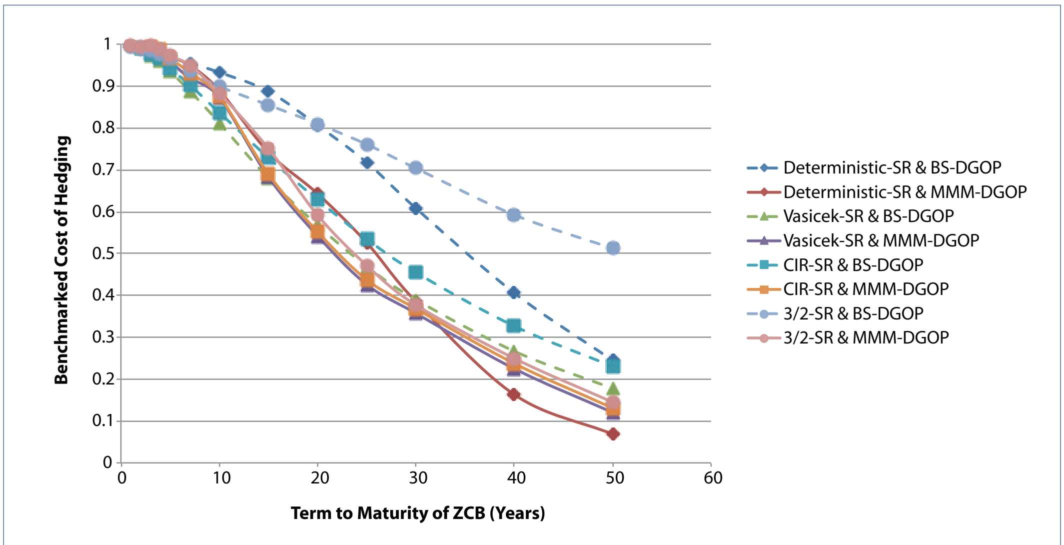

11. HEDGING SWAPTIONS

Swaptions are interest rate derivatives which protect the owner of such an asset against a rise or fall in swap rates and are therefore used by many pension funds and life insurers seeking to hedge their exposure to interest rates. For example, low interest rates may be bad for some life insurers because it becomes expensive to invest in fixed income products, which match their bond-like liabilities.

In this section we describe the hedging strategy used to replicate the payoff of a 3%-strike payer swaption at expiry having an underlying semi-annual 10-year swap and unit notional amount. As in Section 1 we calculate the cost of purchasing a hedge portfolio at the outset of hedging the swaption which ensures that the swaption payoff is affordable by the hedge strategy with high probability. We compute the costs of hedging swaptions across all market models and all swaption times to expiry and identify the best performing models. We deliberately focus on swaptions that have an underlying 10-year swap rate because our US data set contains ten-year swap rates and therefore allows us to calculate a payoff at expiration of the swaption.

We describe how we price swaptions in this paper. The payoff at time T of a payer swaption with unit notional and strike rate R is:

(35) $$H_T = \sum_{i=1}^{n} P(T,T_i)(T_i - T_{i-1})(SW_T - R)^+$$

where $T_1, \dots, T_n$ are the payment times of the underlying swap, SW_T is the corresponding swap rate at time T and we take $T_0 = T$. This payoff is the same as that of a put option on a coupon bond with coupon rate R and having strike price equal to one, that is:

(36) $$H_T = \left(1 - \sum_{i=1}^{n}(R(T_i - T_{i-1}) + 1_{i=n})P(T,T_i)\right)^+$$

see for example Hull [1997]. Here $1_{i=n}$ denotes the indicator function which equals one if $i = n$ and zero otherwise. Applying (8) to this payoff for the deterministic short rate and Black-Scholes discounted GOP market model, the real world swaption pricing formula simplifies to the intrinsic value of the swaption, namely:

(37) $$V_t^{\delta_{H_T}} = \big(1 - \sum_{i=1}^{n}(R(T_i - T_{i-1}) + 1_{i=n})\exp(-\int_t^{T_i} r_s ds)\big)^+$$

For a mean reverting Gaussian interest rate model Jamshidian 1989 proves that the price of a coupon bond option with strike price K, and corresponding strike rate R, is equal to the sum of options on constituent zero coupon bonds each having its strike price calculated from the common strike rate R. The derivation of the formula relies on the observation that the monotonicity of the zero coupon bond price as a function of the short rate implies that the exercise short rate of the portfolio of ZCBs is the same as each of the exercise short rates of the options on the component ZCBs (see for example Hull & White 1990). Jamshidian's formula suffices for the Vasicek short rate and Black-Scholes discounted GOP model and in this instance we have the real world swaption pricing formula:

(38) $$\begin{aligned} V_t^{\delta_{H_T}} = \sum_{i=1}^{n} \big(R(T_i - T_{i-1}) + 1_{i=n}\big)\big(&- A(t,T_i)\exp(-r_t B(t,T_i))N(-d_i^{(1)}) \\ &+ A(t,T)\exp(-r_t B(t,T))K_i N(-d_i^{(2)})\big) \end{aligned}$$

where A and B are given in (13) and (12) and $d_i^{(1)}$ and $d_i^{(2)}$ are given by:

(39) $$d_i^{(1)} = \frac{1}{\sigma_i}\log\left(\frac{A(t,T_i)\exp(-r_t B(t,T_i))}{A(t,T)\exp(-r_t B(t,T))K_i}\right) + \frac{1}{2}\sigma_i$$

(40) $$d_i^{(2)} = \frac{1}{\sigma_i}\log\left(\frac{A(t,T_i)\exp(-r_t B(t,T_i))}{A(t,T)\exp(-r_t B(t,T))K_i}\right) - \frac{1}{2}\sigma_i$$

with σ_i given by:

(41) $$\sigma_i = \sigma B(T,T_i)\sqrt{\frac{1}{2\kappa}(1 - \exp(-2\kappa(T-t)))}$$

and K_i given by:

$$K_i = A(T,T_i)\exp(-xB(T,T_i)) \tag{42}$$

Here x is the solution to the equation:

$$1 = \sum_{i=1}^{n}(R(T_i - T_{i-1}) + 1_{i=n})A(T,T_i)\exp(-xB(T,T_i)) \tag{43}$$

Also, Jamshidian's method can be adapted to the deterministic short rate and MMM discounted GOP model giving the real world swaption pricing formula:

$$\begin{aligned} V_t^{\delta_{H_T}} = \sum_{i=1}^{n}(R(T_i - T_{i-1}) + 1_{i=n})\Bigg(& -\frac{B_t}{B_{T_i}}(\chi^2_{0,\lambda}(u_i^*) - \exp(-\lambda/2)) \\ & + \frac{B_t}{B_{T_i}}\left(\exp\left(-\frac{\tau_i}{1+2\tau_i}\lambda\right)\chi^2_{0,\lambda/(1+2\tau_i)}((1+2\tau_i)u_i^*) - \exp(-\lambda/2)\right) \\ & + K_i\frac{B_t}{B_T}(\chi^2_{0,\lambda}(u_i^*) - \exp(-\lambda/2))\Bigg) \end{aligned} \tag{44}$$

where:

$$u_i^* = \begin{cases} 2\frac{\varphi_{T_i}-\varphi_T}{\varphi_{T_i}-\varphi_t}\log\frac{1}{1-K_iB_{T_i}/B_T} & \text{if} \quad 1 > K_iB_{T_i}/B_T \\ \infty \quad \text{otherwise} \end{cases} \tag{45}$$

$$\lambda = \frac{\bar{S}_t^{\delta_*}}{\varphi_T - \varphi_t}$$

$$\tau_i = \frac{1}{2}\frac{\varphi_T - \varphi_t}{\varphi_{T_i} - \varphi_T}$$

and K_i given by:

$$K_i = 1 - \exp(-x\tau_i) \tag{46}$$

Here x is the solution to the equation:

$$1 = \sum_{i=1}^{n}(R(T_i - T_{i-1}) + 1_{i=n})(1 - \exp(-x\tau_i)) \tag{47}$$

We have used the notation $\chi^2_{\nu,\lambda}(x)$ to denote the cumulative distribution function of a non-centrally distributed random variable having non-centrality parameter λ and ν degrees of freedom.

However, for the Vasicek short rate and MMM discounted GOP model we resort to the following theorem.

Theorem

Suppose that the short rate r_t obeys Vasicek's SDE (2) and suppose the short rate and the discounted GOP $\bar{S}_t^{\delta_*}$ are independent. Then the real world price of a coupon bond put option is given by:

$$V_t^{\delta_{H_T}} = \int_0^\infty \frac{\lambda}{u}V_t^{\delta_{H_T}}(u)f_{\chi^2_{4,\lambda}}(u)\,du \tag{48}$$

where:

$$\begin{aligned} V_t^{\delta_{H_T}}(u) = \sum_{i=1}^{n}\Bigg(& \left(R(T_i - T_{i-1}) + 1_{i=n}\right) \times (1 - \exp(\tau_i u)) \\ & \times \big(-A(t,T_i)\exp(-r_tB(t,T_i))N(-d_i^{(1)}(u)) \\ & + A(t,T)\exp(-r_tB(t,T))K_i(u)N(-d_i^{(2)}(u))\big)\Bigg) \end{aligned} \tag{49}$$

where A and B are given in (13) *and* (12) *and* $d_i^{(1)}(u)$ *and* $d_i^{(2)}(u)$ *are given by:*

(50) $$d_i^{(1)}(u) = \frac{1}{\sigma_i} \log\left(\frac{A(t,T_i)\exp(-r_t B(t,T_i))}{A(t,T)\exp(-r_t B(t,T))K_i(u)}\right) + \frac{1}{2}\sigma_i$$

(51) $$d_i^{(2)}(u) = \frac{1}{\sigma_i} \log\left(\frac{A(t,T_i)\exp(-r_t B(t,T_i))}{A(t,T)\exp(-r_t B(t,T))K_i(u)}\right) - \frac{1}{2}\sigma_i$$

with σ_i *given by:*

(52) $$\sigma_i = \sigma B(T,T_i)\sqrt{\frac{1}{2\kappa}(1-\exp(-2\kappa(T-t)))}$$

and $K_i(u)$ given by:

(53) $$K_i(u) = A(T,T_i)\exp(-x_u B(T,T_i))$$

Here x_u is the solution to the equation:

(54) $$1 = \sum_{i=1}^{n}(R(T_i - T_{i-1}) + 1_{i=n})(1-\exp(\tau_i u))A(T,T_i)\exp(-x_u B(T,T_i))$$

Proof. We employ the real world pricing formula (8) to the payoff H_T given in (36). The real world price of the bond put option is:

(55) $$\begin{aligned} V_t^{\delta_{H_T}} &= E_t\left(\frac{S_t^{\delta_*}}{S_T^{\delta_*}} H_T\right) \\ &= E_t\left(\frac{\bar{S}_t^{\delta_*}}{\bar{S}_T^{\delta_*}} \frac{B_t}{B_T} H_T\right) \\ &= E_t\left(\frac{\bar{S}_t^{\delta_*}}{\bar{S}_T^{\delta_*}} E_t\left(\frac{B_t}{B_T} H_T \middle| \bar{S}_T^{\delta_*}\right)\right) \end{aligned}$$

Now let $U = \bar{S}_T^{\delta_*}/(\varphi_T - \varphi_t)$, which is a non-central chi-squared random variable having four degrees of freedom and non-centrality parameter λ (see Platen & Heath 2006). We can write the inner expectation above as:

(56) $$\begin{aligned} V_t^{\delta_{H_T}}(u) &= E_t\left(\frac{B_t}{B_T} H_T \middle| \bar{S}_T^{\delta_*} = (\varphi_T - \varphi_t)u\right) \\ &= E_t\left(\frac{B_t}{B_T} H_T \middle| U = u\right) \\ &= E_t\left(\frac{B_t}{B_T}\left(1 - \sum_{i=1}^{n}(R(T_i - T_{i-1}) + 1_{i=n})(1-\exp(\tau_i u))E_T\left(\frac{B_T}{B_{T_i}}\right)\right)^+\right) \end{aligned}$$

which we see is the value of a put option on a portfolio of zero coupon bonds. By analogy to (38) we have:

(57) $$\begin{aligned} V_t^{\delta_{H_T}}(u) = \sum_{i=1}^{n} \Big(&\big(R(T_i - T_{i-1}) + 1_{i=n}\big) \times (1-\exp(\tau_i u)) \\ &\times \big(-A(t,T_i)\exp(-r_t B(t,T_i))N(-d_i^{(1)}(u)) \\ &+ A(t,T)\exp(-r_t B(t,T))K_i(u)N(-d_i^{(2)}(u))\big)\Big) \end{aligned}$$

where A and B are given in (13) and (12) and $d_i^{(1)}(u)$ and $d_i^{(2)}(u)$ are given by:

(58) $$d_i^{(1)}(u) = \frac{1}{\sigma_i} \log\left(\frac{A(t,T_i)\exp(-r_t B(t,T_i))}{A(t,T)\exp(-r_t B(t,T))K_i(u)}\right) + \frac{1}{2}\sigma_i$$

(59) $$d_i^{(2)}(u) = \frac{1}{\sigma_i} \log\left(\frac{A(t,T_i)\exp(-r_t B(t,T_i))}{A(t,T)\exp(-r_t B(t,T))K_i(u)}\right) - \frac{1}{2}\sigma_i$$

with σ_i given by:

(60) $$\sigma_i = \sigma B(T,T_i)\sqrt{\frac{1}{2\kappa}(1-\exp(-2\kappa(T-t)))}$$

and $K_i(u)$ given by:

(61) $$K_i(u) = A(T,T_i)\exp(-x_u B(T,T_i))$$

Here x_u is the solution to the equation:

(62) $$1 = \sum_{i=1}^{n}(R(T_i - T_{i-1}) + 1_{i=n})(1-\exp(\tau_i u))A(T,T_i)\exp(-x_u B(T,T_i))$$

Therefore:

(63) $$\begin{aligned} V_t^{\delta_{H_T}} &= E_t\left(\frac{\bar{S}_t^{\delta_*}}{\bar{S}_T^{\delta_*}} V_t^{\delta_{H_T}}\left(\frac{\bar{S}_T^{\delta_*}}{\varphi_T - \varphi_t}\right)\right) \\ &= \int_0^\infty \frac{\lambda}{u} V_t^{\delta_{H_T}}(u) f_{\chi^2_{4,\lambda}}(u)\,du \end{aligned}$$

as required.

Table 11: Percentile costs of hedging swaptions under a deterministic short rate and Black-Scholes discounted GOP based on US data 1871–2012.

Term to expiry of swaption	99th percentile	95th percentile	90th percentile	85th percentile	80th percentile
1Y	0.48611	0.38003	0.31506	0.27528	0.20821
2Y	0.44005	0.33579	0.29116	0.25279	0.20378
3Y	0.4103	0.32175	0.27169	0.2263	0.19394
4Y	0.38279	0.29708	0.23888	0.19915	0.18646
5Y	0.36652	0.26234	0.22598	0.20163	0.18189
7Y	0.306	0.22095	0.20673	0.18653	0.17151
10Y	0.25529	0.21632	0.19987	0.17504	0.15582
15Y	0.24448	0.22389	0.20212	0.17748	0.14858
20Y	0.24321	0.22801	0.21065	0.16979	0.14934
25Y	0.24677	0.22338	0.21134	0.18003	0.15261
30Y	0.23978	0.22031	0.2028	0.17156	0.14505
40Y	0.18559	0.17428	0.17007	0.14754	0.12871
50Y	0.12804	0.11657	0.10833	0.1072	0.10012

Table 12: Percentile costs of hedging swaptions under a deterministic short rate and MMM discounted GOP based on US data 1871–2012.

Term to expiry of swaption	99th percentile	95th percentile	90th percentile	85th percentile	80th percentile
1Y	0.47544	0.37312	0.31076	0.25785	0.21711
2Y	0.42743	0.34451	0.29103	0.24248	0.19133
3Y	0.3962	0.31381	0.27224	0.22923	0.18142
4Y	0.37955	0.29242	0.25508	0.20019	0.18678
5Y	0.3588	0.26402	0.22649	0.19589	0.18186
7Y	0.31133	0.22974	0.2077	0.19088	0.17649
10Y	0.26536	0.22559	0.20246	0.19369	0.18583
15Y	0.23231	0.22473	0.21262	0.20827	0.18639
20Y	0.19473	0.18542	0.1812	0.17025	0.15873
25Y	0.15351	0.14079	0.12989	0.12479	0.11982
30Y	0.13853	0.11254	0.097728	0.084924	0.077339
40Y	0.084456	0.061148	0.046843	0.04057	0.036779
50Y	0.031698	0.025796	0.024644	0.02379	0.023148

12. COSTS OF HEDGING SWAPTIONS UNDER A DETERMINISTIC SHORT RATE

We present the costs of hedging swaptions under market models having a deterministic short rate.

In Table 11 the percentile benchmarked costs of hedging swaptions of various terms to expiry and percentiles are shown for the deterministic short rate and Black-Scholes discounted GOP model.

In Table 12 the percentile benchmarked costs of hedging swaptions of various terms to expiry and percentiles are shown for the deterministic short rate and MMM discounted GOP model. For hedging swaptions with terms to expiry at or shorter than 15 years, the BS discounted GOP model and the MMM discounted GOP model perform similarly. Beyond swaption terms to expiry of 15 years the MMM discounted GOP model significantly outperforms the BS discounted GOP model. For example, hedging a 50Y swaption at the 99th percentile incurs a cost of 0.031698 under the MMM discounted GOP model which is roughly a quarter of the corresponding cost of 0.12804 under the BS discounted GOP model.

13. COSTS OF HEDGING SWAPTIONS UNDER A VASICEK SHORT RATE

We present the costs of hedging swaptions under Vasicek short rate models.

In Table 13 the percentile benchmarked costs of hedging swaptions of various terms to expiry and percentiles are shown for the Vasicek short rate and Black- Scholes discounted GOP model.

In Table 14 the percentile benchmarked costs of hedging swaptions of various terms to expiry and percentiles are shown for the Vasicek short rate and MMM discounted GOP model.

For hedging swaptions with terms to expiry shorter than 7 years, the MMM discounted GOP model and the BS discounted GOP model perform similarly. However, at or beyond swaption terms to expiry of 7 year the MMM discounted GOP model gives lower costs of hedging that given by the Black-Scholes discounted GOP model.

Table 13: Percentile costs of hedging swaptions under a Vasicek short rate and Black-Scholes discounted GOP based on US data 1871–2012.

Term to expiry of swaption	99th percentile	95th percentile	90th percentile	85th percentile	80th percentile
1Y	0.45536	0.35855	0.30709	0.24918	0.21407
2Y	0.39077	0.3202	0.27933	0.23469	0.19518
3Y	0.34031	0.29422	0.25338	0.22579	0.16604
4Y	0.28045	0.25812	0.22965	0.21069	0.15451
5Y	0.2584	0.24	0.20927	0.18904	0.13578
7Y	0.22588	0.19141	0.17426	0.15609	0.12294
10Y	0.19525	0.15656	0.12256	0.11395	0.098862
15Y	0.1472	0.12391	0.095036	0.074855	0.06858
20Y	0.10367	0.09483	0.075293	0.061024	0.053613
25Y	0.071685	0.06604	0.058912	0.049244	0.045019
30Y	0.0502	0.045495	0.043199	0.040535	0.038422
40Y	0.028523	0.027775	0.026965	0.025942	0.025264
50Y	0.018271	0.017919	0.017433	0.016915	0.016747

Table 14: Percentile costs of hedging swaptions under a Vasicek short rate and MMM discounted GOP based on US data 1871–2012.

Term to expiry of swaption	99th percentile	95th percentile	90th percentile	85th percentile	80th percentile
1Y	0.45291	0.34929	0.29782	0.26397	0.20662
2Y	0.41355	0.32358	0.27757	0.24985	0.18266
3Y	0.3485	0.28436	0.24871	0.21766	0.16344
4Y	0.31349	0.24955	0.22659	0.19592	0.15913
5Y	0.26693	0.23529	0.20533	0.17569	0.14353
7Y	0.19943	0.18343	0.17188	0.15911	0.11132
10Y	0.15752	0.13175	0.12518	0.1185	0.094938
15Y	0.10778	0.10203	0.084549	0.070733	0.060173
20Y	0.085821	0.076264	0.068102	0.052175	0.045619
25Y	0.060256	0.057678	0.049517	0.04639	0.037491
30Y	0.043094	0.041715	0.038363	0.034956	0.032305
40Y	0.024757	0.024145	0.0237	0.023144	0.022709
50Y	0.017216	0.01655	0.016265	0.016035	0.014939

14. CONCLUSIONS ON THE SWAPTION HEDGING PERFORMANCES

In Figure 2 the percentile costs of hedging swaptions of varying terms to expiry are graphed. Each model for which the discounted GOP is modelled by the MMM has cheaper costs of hedging long-dated swaptions. In particular, we find that the Vasicek short rate and MMM discounted GOP model provides the cheapest hedging strategy for long-dated swaptions. The market model composed of the Vasicek short rate and MMM discounted GOP was also the best model for hedging zero coupon bonds in Section 10.

Figure 2: Percentile costs of hedging swaptions of varying terms to expiry.

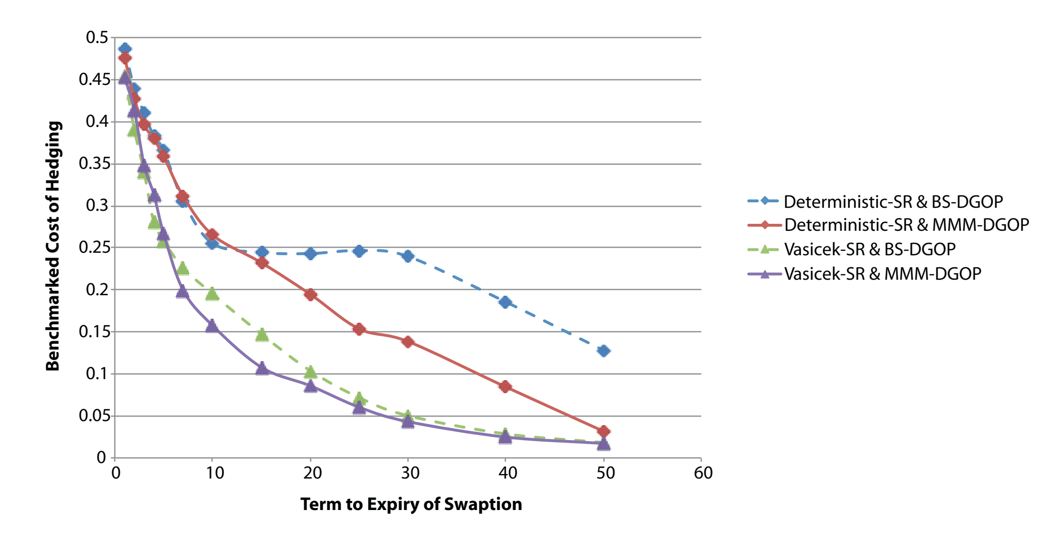

Bibliography

Ahn, D.H., & Gao, B. (1999). A parametric nonlinear model of term structure dynamics. *Review of Financial Studies*, 12:721-762.

Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81:637-654.

Bruti-Liberati, N., & Platen, E. (2010). N*umerical Solution of Stochastic Differential Equations with Jumps in Finance*. Springer.

Cox, J. C., Ingersoll, J. E., & Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica*, 53:385-407.

Hull, J. C. (1997). *Options, Futures, and Other Derivatives*. Prentice-Hall, 3rd edition.

Hull, J. C., & White, A. (1990). Pricing interest rate derivative securities. *Review of Financial Studies*, 3:573-592.

Jamshidian, F. (1989). An exact bond option formula. *Journal of Finance*, 44(1): 205-209.

Platen, E. (2001). A minimal financial market model. *In Trends in Mathematics*, pages 293-301. Birkhäuser.

Platen, E. (2006). On the pricing and hedging of long dated zero coupon bonds. Research Paper 185, University of Technology, Sydney.

Platen, E., & Heath, D. (2006). *A Benchmark Approach to Quantitative Finance*. Springer Finance.

Shiller, R. (1989). *Market Volatility*. MIT Press, Cambridge, Massachusetts.

Dr Kevin Fergusson FIAA
kevin.fergusson@curtin.edu.au
Dr Kevin Fergusson is a Senior Lecturer in Actuarial Science at Curtin University in Perth, WA and has previously worked in the finance industry as an actuarial analyst in superannuation and a quantitative investment manager.

Prof Eckhard Platen
eckhard.platen@uts.edu.au
Professor Eckhard Platen occupies the Chair in Quantitative Finance at the University of Technology, Sydney, bridging the Finance Discipline Group and the School of Mathematical Sciences. He has previously held several leading academic appointments in Berlin and Canberra, co-authored five successful books at Springer Verlag and authored more than 150 research papers in quantitative finance and applicable mathematics. He is currently the President of the Bachelier Finance Society.

Catastrophic mortality bonds: analysing basis risk and hedge effectiveness

ALEX HUYNH, BRIDGET BROWNE AND AARON BRUHN

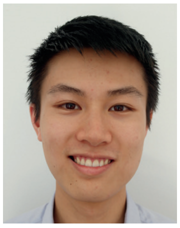

Alex Huynh

Bridget Browne

Aaron Bruhn

ABSTRACT

Life insurers are exposed to catastrophic mortality risk. Catastrophic mortality bonds are a recent market innovation that provide an alternative risk management tool to address this risk. However there is little in the way of published studies that examines their effectiveness, given that they are subject to basis risk arising from the use of country-level general population mortality in their construction. By constructing a typical mortality risk portfolio and calibrating a bond for this portfolio, the hedge effectiveness of the instrument is analysed under a wide variety of circumstances. We find that, on a stand-alone basis, hedge effectiveness may be too low to be acceptable to small to medium insurers. However, effectiveness of the bond increases when used in combination with surplus reinsurance and/or when pooling is used to increase portfolio size.

KEYWORDS

life insurance, mortality risk management, insurance securitisation, catastrophic mortality bonds, basis risk, hedge effectiveness

1 INTRODUCTION

Life insurers and reinsurers are exposed to the risk of future mortality uncertainty. Catastrophic mortality events pose a significant threat to the life insurance industry as they cause a sudden increase in mortality over a short period of time, which may lead to a substantial rise in claims and the potential for severe adverse financial consequences, such as breaches in regulatory solvency and capital requirements (Cox & Hu 2004).

Although there are a range of catastrophic mortality events that may impact the life insurance industry, an influenza pandemic is considered the most serious threat. The exposure to catastrophic mortality events such as influenza pandemics has been difficult for life insurers and reinsurers to manage since the probability of such events occurring in any year is low while the potential for devastating losses is high. Catastrophic mortality bonds are a recent capital market innovation that provide an alternative risk management tool to hedge against catastrophic mortality events. In contrast to the inherent credit risk associated with traditional reinsurance, they bear essentially no credit risk for sponsors (Bagus 2007). However, catastrophic mortality bonds are not indemnity based, as the payoff trigger is based on a specified mortality index, which is calculated as a weighted average of general population mortality rates (Cowley & Cummins 2005). Consequently, the issue of basis risk arises, resulting in imperfect hedge effectiveness as the possibility exists for gains or losses in the hedged position. In particular, the sponsor is concerned that the bond payoff will be inadequate to cover the actual loss suffered (Coughlan et al. 2011).

This article quantifies the basis risk and hedge effectiveness of catastrophic mortality bonds in order to explore the level of coverage they provide for sponsors. It examines the use of catastrophic mortality bonds in hedging against additional claims arising from an influenza pandemic using an individual fully underwritten yearly renewable term (YRT) insurance portfolio as a practical example. It does not attempt to develop a definitive catastrophic mortality model, as this is not required for our purposes here. Interested readers may use this as a guide for application to their own specific portfolios.

Overall, we find that there is significant variation in the basis risk and hedge effectiveness of catastrophic mortality bonds. The findings suggest that catastrophic mortality bonds are a viable alternative risk management tool for large portfolio sizes, for portfolios where the distribution of exposed sums insured is less spread and where the life insurer's underlying exposure remains relatively stable. Hence the pooling of small to medium portfolios and/or the use of surplus reinsurance to homogenise the distribution of sums insured may be effective risk management strategies to adopt concurrently with implementation of a catastrophic mortality bond.

The remainder of this article is structured as follows. Section 2 briefly outlines the epidemiological characteristics of influenza pandemics in order to set the scene for the calibration of a bond issuance and discusses the life insurer's management of catastrophic mortality risk arising from an influenza pandemic. Section 3 provides an overview of the catastrophic mortality bond market, the key features of the bonds and concepts of basis risk and hedge effectiveness. Section 4 describes the methodology adopted for the analysis. Section 5 reports the results obtained and summarises key findings. Section 6 concludes.

2 INFLUENZA PANDEMICS

2.1 Epidemiological characteristics

While seasonal influenza epidemics usually occur during the autumn and winter months in temperate regions and all year round for tropical and sub-tropical regions, the emergence of influenza pandemics is not constrained by season (Nguyen-Van-Tam & Hampson 2003) and it is reasonable to believe that influenza pandemics may appear at any time during the year.

In contrast to seasonal influenza epidemics that occur annually, influenza pandemics are rare and unpredictable events, which have occurred irregularly throughout history. To date, there is no identified chronological pattern that would allow us to predict the occurrence of future influenza pandemics (Potter 2001). Influenza pandemics have been characterised by multiple waves of infection with varying impact occurring over two calendar years. A variety of patterns have been observed regarding duration and severity, which will be incorporated in later modelling.

This paper does not attempt to develop a sophisticated or micro-level approach to modelling pandemic risk, as this is not required to illustrate the issue of basis risk.

2.1.1 Excess mortality rate

Evidence suggests that influenza pandemics may cause considerably higher excess mortality (defined as the difference between the observed mortality rate and the expected baseline mortality rate in the absence of an influenza pandemic (Simonsen et al. 1997)), but this impact is difficult to quantify because influenza may not be listed as a cause on the death certificate for many influenza related deaths (Woolnough et al. 2007). Consequently, the all-cause excess mortality and influenza- and pneumonia-specific excess mortality can be considered as the upper and lower bounds of mortality attributed to an influenza pandemic, respectively. The excess mortality rate has varied significantly among the influenza pandemics of the last 100 years, as illustrated in Table 1.

Table 1: Estimated excess mortality rates for the influenza pandemics of the 20th and 21st century.

Name	Global excess mortality rate (per 1,000) [a]	US excess mortality rate (per 1,000) [a]
1918–1919 Spanish Flu	27.60–55.20	4.81–6.50
1957–1958 Asian Flu	0.34–0.69	0.38–0.46
1968–1969 Hong Kong Flu	0.28	0.14–0.17
2009–2010 H1N1 Flu	Not available	0.02–0.14

Sources: Dauer & Serfling (1961); Glezen (1996); Simonsen et al. (1998); US Census Bureau (2000, 2011a, 2011b); US Department of Health and Human Services (2011); United Nations (1999); Viboud et al., (2010); World Health Organisation (2005)
[a] The excess mortality rate is calculated as the number of excess deaths divided by the average of the population over the influenza pandemic.

Death rates during pandemic events are unlikely to satisfy the standard (and necessarily simplified) assumption of independence between lives. However, we will make this assumption in later modelling for simplicity, the implication being that the results may be somewhat understated.

2.1.2 Age-specific distribution of excess mortality rates

The age-specific distribution of excess mortality rates for seasonal influenza epidemics typically has a U-shaped curve, representing high mortality among infants and the elderly with comparatively low mortality rates at ages in between (Nguyen-Van-Tam & Hampson 2003). On the other hand, the age-specific distribution of excess mortality rates for influenza pandemics has tended to affect a higher proportion of persons under 65 years of age than seasonal influenza. This is often attributed to the partial immunity that many persons over 65 years of age may have retained from exposure to similar influenza infections as children or young adults (Nguyen-Van-Tam & Hampson 2003). The age-specific distribution of excess mortality rates for the last four influenza pandemics have exhibited either U, \/\, or /\ shapes and have been similar for both genders. Further details on the level and characteristics of these distributions may be found in Huynh et al. (2013).

2.2 Catastrophic mortality risk management

2.2.1 Risk identification

For most life insurers, death benefit products constitute the majority of their risk business and, as a result, they are likely to experience a significant loss from an influenza pandemic. This has the potential to severely impact the life insurer's results and may lead to breaches of solvency requirements.

2.2.2 Risk measurement

Life insurers assess the potential impact of influenza pandemics with risk modelling and scenario testing. Internal risk models are commonly used to assess a full range of (perceived) risks and to take into account dependencies between different risks and exposures, which can be complex. Since influenza pandemics are rare events, there is scarce data to calibrate a number of uncertain parameters required in the models. This is typically dealt with through sensitivity testing (Baumgart et al. 2007).

Several studies have examined the potential impact by estimating the additional cost based on a range of deterministic scenarios derived from historical influenza pandemics. It is apparent from Table 2 that a wide range of outcomes are considered possible. In general, the studies conclude that the life insurance industry can absorb the impact of a severe pandemic, but will incur significant decreases in profit and capital. It is also noted that the life reinsurance industry will be more heavily impacted, since reinsurance is essentially pure mortality risk business, and that the advantage conferred by reinsurers' geographical diversification ceases to apply in the event of a pandemic (Dreyer et al. 2007 APRA 2007).

2.2.3 Risk management

Retention of catastrophic mortality risk is possible, but is unlikely to be economically efficient. Clearly, geographic diversification is not as effective as with other, localised, catastrophic mortality events since an influenza pandemic is likely to affect multiple geographical regions around the world. Diversification across lines of businesses is also somewhat limited, because health and general insurance business may also be affected. On the other hand, annuity business may provide a natural hedge as the value of protection and annuity liabilities move in the opposite direction in response to a change in mortality (Cox & Lin 2007). This will, however, depend on the age-specific distribution of excess mortality rates, as life insurers primarily write protection policies to younger age groups and sell annuity policies to older age groups.

Life insurers can transfer mortality risk through reinsurance. However, this exposes the insurer to credit risk because the reinsurer may develop solvency issues

Table 2: Summary of assumptions and results from studies examining the potential impact of an influenza pandemic on the life insurance industry.

Author(s)	Country	Severity [a]	General population excess mortality rate (per 1,000)	Modelled age-specific distribution of excess mortality rate	Excess mortality rate ratio of insured versus general population (%)[b]	Influenza pandemic duration (Years)	Results: additional gross claims (AGC) or additional net claims (ANC)
APRA (2007)	Australia	Severe	1.0	Flat	100	1	AGC: AUD 1.2 billion
Dreyer, Kritzinger & Decker (2007)	South Africa	Mild Moderate Severe	0.40 1.40 20.0	‘W’ ‘W’ ‘W’	Group life: 70 Individual life: 40*	1 1 1	AGC: ZAR 0.753 billion AGC: ZAR 2.7 billion AGC: ZAR 37.6 billion
Stracke (2007)	Germany	Severe	6.4	‘W’	100	1	ANC: EUR 5.1 billion
Toole (2007)	USA	Moderate Severe	0.70 6.5	‘U’ ‘\/\’	57.1 76.9	1 1	ANC: USD 2.8 billion ANC: USD 64.3 billion
Weisbart (2006)	USA	Moderate Severe	1.07 4.81	‘_/’ ‘U’	100 100	1 1	AGC: USD 31 billion AGC: USD 133 billion

[a] Severity ratings are specified by each of the authors of the study.
[b] The excess mortality rate ratio of insured versus general population reflects the potential for better mortality experience in the insured population relative to the general population in an influenza pandemic scenario.
* The excess mortality rate ratio of insured versus general population for Dreyer, Kritzinger & Decker (2007) is the same across all three scenarios, but is differentiated into group life and individual life products.

due to a pandemic event causing them to default on their obligations or be slow to pay reinsurance claims (Baumgart et al. 2007). An alternative is a catastrophic mortality bond, which essentially eliminates credit risk when well designed. These instruments offer several advantages and disadvantages compared to reinsurance, which are discussed in Huynh et al. (2013) along with an introduction to the market and key features of these instruments. A key disadvantage is the presence of basis risk, which is discussed in the following section.

3 CATASTROPHIC MORTALITY BONDS – BASIS RISK AND HEDGE EFFECTIVENESS

Basis risk arises whenever there are differences between an underlying hedged portfolio and the associated hedging instrument. Its presence implies imperfect hedge effectiveness because there is a possibility of gains or losses in the hedged position. This does not necessarily invalidate the case for hedging because basis risk can be minimised by appropriately structuring and calibrating the hedging instrument to ensure high hedge effectiveness. If the basis risk is small relative to the risk of the initial unhedged position then it is possible for the hedging strategy to be beneficial (Coughlan et al. 2011).

The issue of basis risk has been examined for several index-based insurance linked securities (ILS). In non-life, industry loss warranties, catastrophic loss index securities and catastrophe insurance linked contracts have been examined (see, for example, Cummins et al. 2004; Harrington & Niehaus 1999 Zeng, 2000). In life, the extant literature has primarily focused on longevity linked securities (see, for example, Coughlan et al. 2011; Cairns et al. 2011; Ngai & Sherris 2011). To the authors' knowledge, there had been no published literature on the analysis of basis risk and hedge effectiveness for catastrophic mortality bonds at the time of writing.

In the context of catastrophic mortality bonds, basis risk could arise from differences between general and insured populations due to initial or emerging mismatches in age, gender, geographical location and socioeconomic class (Coughlan et al. 2011).[1] As age, gender, country and financial exposure in the form of sum assured can be calibrated to match that of the life insurer's initial underlying exposure, one important determinant of basis risk is therefore associated with mismatches of socioeconomic class (Richards & Jones, 2004).

Some historical evidence from an 1889 pandemic (Mead 1919) and the 1918–19 pandemic in Craig and Dublin (1919) and Little (1919) suggests that the impact of underwriting and economic self-selection will continue to result in lighter mortality experience in the insured

1 This is additional to the issue of differing overall mortality experience between general and insured populations. Insured mortality may be significantly lower than that of the general population due to the impact of underwriting and economic self-selection, with differences dependent on age, gender, smoking class, policy duration and underwriting type (Toole 2007)

population in the event of an influenza pandemic, as compared with the general population. In addition, a study on the 1957–1958 and 1968–1969 influenza pandemics observes approximately 12% lower excess death rates in standard ordinary policyholders compared with age- and gender-matched general population (Woolnough et al. 2007), consistent with other findings such as Metropolitan Life Insurance Company (1976).

However, basis risk remains even if similar characteristics are shared or differences are calibrated for, simply because the two populations are not the same people.

4 METHODOLOGY

4.1 Overview of framework

The framework for assessing basis risk and hedge effectiveness is adapted from Coughlan et al. (2011) and is shown in Figure 1. The life insurer's claims are determined by the insured population mortality rates while the catastrophic mortality bond payoff is determined by the general population mortality rates. The evaluation of hedge effectiveness uses simulations of the forecast mortality rates to calculate these cash flows, which are used to calculate the hedge effectiveness.

Figure 1: Framework for assessing basis risk and hedge effectiveness.

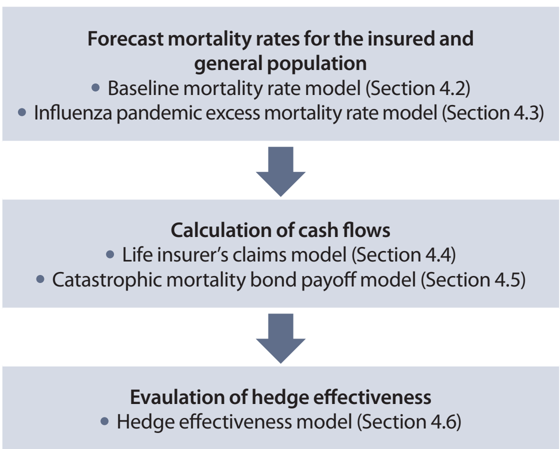

Source: Modified from Coughlan et al. (2011)

4.2 Baseline mortality rate model

The baseline mortality model forecasts the baseline mortality rate for the general and insured population. The baseline mortality rate is defined as the future annual mortality rates assuming that no catastrophic mortality event occurs. In this paper, Australian mortality is used, however the methodology would be the same for any country or countries where an insurer may have exposure.

4.2.1 Life tables and mortality improvements

The Australian Bureau of Statistics (ABS) 2007–2009 life tables are the basis for the Australian general population mortality. The mortality of the Australian insured population is based on the IA95-97 life tables published by the Institute of Actuaries of Australia.

Both life tables are assumed to improve by the Australian Government Actuary (AGA) 25-year mortality improvement trend to 2010, to establish the starting point for the bond. For projecting beyond 2010, a divergence in the rate of change in mortality rates could have an impact on hedge effectiveness. Therefore, three possible mortality improvement trends are considered for both general and insured populations: the AGA 25 year mortality improvement trends, the 100-year mortality improvement trends, and no mortality improvement.

4.3 Influenza pandemic excess mortality rate model

This component forecasts the excess mortality rate for the general and insured population in the event that an influenza pandemic occurs. The excess mortality rate is modelled explicitly and not decomposed into the clinical attack rate and case fatality rate, and is modelled only as age-specific and not gender-specific.

It is uncertain whether increases in mortality caused by a historical influenza pandemic should be treated as additive (i.e. absolute) or multiplicative (i.e. relative) to baseline mortality rates for the purpose of estimating future impact of a 'similar' pandemic.[2] In accordance with existing studies, we have modelled the pandemic as additive.

There are two broad sources from which to develop the assumption: actual historical influenza pandemic mortality data as described in Section 2.1 and assumptions used by studies examining the potential impact of an influenza pandemic as described in Section 2.2.2. The 20th century influenza pandemic mortality data is difficult to apply to today's situation, due to significant environmental changes including improvements in medical care and technology, establishment of global health monitoring and early warning systems, better communication methods and improved socio-economic conditions (Baumgart et al. 2007). On the other hand, some changes may increase the impact of future influenza pandemics such as a higher percentage of the population at older ages,

2 For example, the Spanish Flu caused an additive increase to baseline mortality rates of approximately 5 per 1,000 or a multiplicative increase of 30% of baseline mortality rates. But because baseline mortality rates have decreased over time, the two approaches produce quite different results when applied to current baseline mortality rates. To continue the above example, applying the additive increase to the Australian standardised baseline mortality rate of 5.7 per 1,000 results in a multiplicative increase of 88% instead of 30%.

increased urban population density and increased human mobility (Faulds & Bridel 2009). There are no explicit adjustments made to account for these changes in this paper given the substantial uncertainty surrounding their impact. In any case, the range of historical influenza pandemic severities already provides a wide range of potential scenarios.

4.3.1 Overall general population excess mortality rate

The range of general population excess mortality rate assumptions used by the influenza pandemic studies is broadly consistent with the range of US excess mortality during the past four influenza pandemics (as shown in Table 1 and Table 2). The 1918–1919 Spanish Flu is generally considered as the upper bound on future influenza pandemic mortality, even though there is evidence to prove that this represent the maximum possible mortality (Murray et al. 2006). Notwithstanding, historical experience suggests that the general population excess mortality rate may vary between 0 and 5 per 1,000.

4.3.2 Age-specific distribution of excess mortality rates

For the purposes of this paper, four shapes of age-specific distribution of excess mortality rates are considered: U, \/\ and W and a flat curve. These draw on the experience discussed in Section 2.1.2 and the modelling presented in Table 2. Each shape is calibrated to the age structure of (in this case) Australia to ensure that the overall excess mortality rate assumed to be observed is consistent with the actual overall excess mortality rate.

4.3.3 Excess mortality rate ratio of insured versus general population

Potential differences in mortality between insured and general populations in the event of a pandemic are unclear, since the effects of underwriting and economic self-selection may cease to apply. As a result of their higher socio-economic status the insured population may have better access to healthcare and be more educated about the impact of influenza, but on the other hand also more likely to engage in international travel and live in high density urban areas. A lack of data exacerbates this uncertainty. Woolnough (2007) suggests an insured to general population mortality ratio of 88% from the 1957–1958 and 1968–1969 influenza pandemics, and as per Table 2 the ratio assumed by influenza pandemic studies ranges from 40% to 100%.

In this paper, the mortality ratio of the insured versus general population is assumed to vary from 40% to 120% to capture the uncertainty regarding this relationship.

4.3.4 Duration of the influenza pandemic and severity of waves each year

Historical evidence from Section 2.1 suggests influenza pandemics may begin at any time during the year and also indicates that influenza pandemics have varying severity of waves over two calendar years. In contrast, influenza pandemic studies assume that the duration is one year, as shown in Table 2. We assume that the influenza pandemic occurs with varying severity of waves each year over one, two or three calendar years.

4.4 Life insurer's claims model

The life insurer's claims model simulates the aggregate annual claims for a typical life insurer's portfolio of individual fully underwritten YRT insurance. This requires assumptions about the number of policies, average sum insured and distribution of sum insured by age and gender.

The aggregate annual simulated claims include sampling risk, which is the risk that the 'realised' mortality is different from the 'intrinsic' mortality due to a small population size. A Bernoulli distribution is used to model the death process of each policyholder where the policyholder dies if a simulated random number between 0 and 1 is less than the policyholder's mortality rate. This is likely to be an understatement of the true pandemic mortality risk, since, as identified in 2.1.1, death rates during pandemics are almost certainly not independent events; however this is acceptable to demonstrate the minimum basis risk that may present.

The portfolio is assumed to comprise 20,000 policies. A relatively small portfolio size is chosen for ease of calculation and because it is more effective in illustrating the relative impacts on hedge effectiveness from varying key parameters. The portfolio composition by age and gender is as shown in Table 3. The smoker status of policyholders was not considered as baseline mortality and influenza pandemic excess mortality rates are not subdivided by smoker status (although this may also contribute to basis risk).

Table 3: Portfolio composition by age and gender.

Age band	Male	Female	Age band	Male	Female
15–24*	0.6%	0.8%	55–64	6.2%	1.5%
25–34	10.6%	9.1%	65–74	0.0%	0.0%
35–44	22.9%	16.7%	75+	0.0%	0.0%
45–54	21.7%	9.9%	Subtotal	62.0%	38.0%

Source: Hayes (2008), as per IA95-97 life tables
* A minimum age of 18 is assumed.

The portfolio composition remains static over the risk period. The average sum insured is assumed to be $365,000, based on a 2007 YRT average sum insured of approximately $330,000 (Palmer 2009). Table 4 provides the gender-specific distributions of sum insured by age.

4.5 Catastrophic mortality bond payoff model

4.5.1 Construction of contingent claim payoff mechanism

The catastrophic mortality bond payoff component models the contingent claim payoff mechanism, as described in the appendix and constructed following Swiss Re Capital Markets (2008).

4.5.2 Calibration of the catastrophic mortality bond

The catastrophic mortality bond must be calibrated according to the life insurer's hedging objective, which for this paper is to protect against significant additional claims arising from a one year influenza pandemic causing a general population excess mortality rate of 1 to 2 per 1,000 with a \/\ age-specific distribution. That is, we do not expect the insurer to want to pay for protection against every additional claim arising, they will only require protection once additional claims reach a financially stressful level.

Assuming an excess mortality rate ratio of insured versus general population of 80%, this equates to an insured population excess mortality rate of 0.8 to 1.6 per 1,000. The life insurer intends to retain the loss from claims caused by all baseline mortality plus excess mortality up to 0.8 per 1,000, referred to as 'retained claims'.

A general population excess mortality rate of 1 to 2 per 1,000 is chosen as this results in attachment and exhaustion points that are consistent with the range observed in the market. The \/\ shape is chosen because this shape has the highest impact on insured lives.

The start of the risk period is assumed to be 1st January 2011, the maturity of the bond is 5 years, and the influenza pandemic is arbitrarily assumed to occur in 2013 and last for one year. We assume that the general and insured population mortality improvement follow the AGA 25 year trend.

For simplicity, there is only one tranche. The attachment and exhaustion point are chosen to be 122.33% and 151.77%, as this is equal to the expected mortality index values with the given mortality assumptions at an excess mortality rate of 1 and 2 per 1,000. The size of age bands is set to 5 years consistent with previous transactions, while the age and gender weightings were fit to the portfolio. The principal amount is $6,466,841, set to the expected aggregate claims between an excess mortality rate of 1 and 2 per 1,000. It should be noted that in reality the characteristics of a bond are also influenced by investor demand.

It is assumed that there is no sampling risk in the general population (in contrast to the life insurer's claims model).

Table 4: Male and female distribution of sum insured by age band.

Male / Female Age-band	Sum insured bands				
	$20k–$150k	$150k–$500k	$500k–$1m	$1m–$3m	$3m–$5m
34 and less	55% / 65%	33% / 26%	10% / 8%	2% / 1%	0% 0%
35–54	40% / 60%	29% / 24%	20% / 12%	9% / 3%	2% / 1%
55 & greater	70% / 80%	20% / 13%	8% / 6%	2% / 1%	0% / 0%

4.6 Hedge effectiveness model

Hedge effectiveness can be defined as the degree of risk reduction of the unhedged exposure for a given risk metric (Cummins et al. 2004; Li & Hardy 2011; Coughlan et al. 2011; Cairns et al. 2011). The unhedged and hedged exposures are defined as the value of the liability and the value of the liability plus the value of the hedging instrument, respectively. In this paper we define the hedge effectiveness (hereafter HE) realised for any specific simulation outcome in a given scenario as follows:

$$HE = 1 - \frac{\text{AC} - \text{RC} - \text{Bond Payoff}}{\text{AC} - \text{RC}}$$

Where:

HE	The hedge effectiveness measure
AC	Actual aggregate claims incurred
RC	Retained claims, i.e. those the insurer was prepared to accept for their own account, caused by an general population excess mortality rate of 1 per 1,000
Bond Payoff	The bond payoff

The numerator is the hedged exposure and the denominator is the unhedged exposure. This means HE will be 100% when the bond payoff is exactly equal to the excess claims incurred (i.e. AC – RC): those against which the bond was intended to protect. Under-hedging occurs when HE is less than 100% indicating that the bond payoff is insufficient to cover the excess claims incurred. Over-hedging, when HE is greater than 100%, occurs when the bond payoff exceeds the excess claims incurred. HE cannot be less than zero as the bond payoff cannot be negative.

Estimated HE is examined at the median as well as the mean. A key measure is the effectiveness of the hedge under more extreme circumstances, hence HE at the 5th percentile, is examined as it indicates the minimum level of coverage with high probability.

We run 10,000 simulations at a general population excess mortality rate of 1 and 1.5 per 1,000 to obtain an empirical distribution of the retained claims and actual aggregate claims, respectively. An experienced excess mortality rate of 1.5 per 1,000 is chosen as it corresponds to the middle of the range of protection targeted. Each simulation generates a value for the actual aggregate claims and retained claims, which are used to calculate a value for the HE. The bond payoff is constant for a given set of assumptions, as it is determined by observation at population level

4.7 Sensitivity analysis

A sensitivity analysis is conducted on the characteristics of the life insurer's portfolio. In performing this analysis, the bond issuance is recalibrated for scenarios (A) 1. and (A) 2. below but not for scenario (A) 3. The changes that are considered include:

- (A) 1. increasing the number of policies in the portfolio to 40, 60, 80 and 100 thousand, produced by replicating the original portfolio
- (A) 2. a flat distribution of sum insured with an equivalent average sum insured
- (A) 3. a portfolio composition with the age of policyholders increased and decreased by 5 years.

Given the inherent uncertainty surrounding future mortality, a sensitivity analysis is also conducted on the mortality assumptions. The changes that are considered include:

- (B) 1. combinations of the AGA 25-year mortality improvement, 100-year mortality improvement, and no mortality improvement for each of the general and insured populations
- (B) 2. an overall general population excess mortality rate (per 1,000) of 1.00, 1.25, 1.75, 2.00 and 2.25
- (B) 3. an age-specific distribution of excess mortality rates of U, W and flat
- (B) 4. an excess insured to general population mortality rate ratio of 40%, 60%, 100% and 120%
- (B) 5. the influenza pandemic occurs in 2011, 2012, 2014, or 2015
- (B) 6. an influenza pandemic duration of 2 and 3 years with different severity of waves each year. The scenarios examined include a 2 year influenza pandemic with an excess mortality rate of 1 and 0.5 per 1,000 in the first and second year, respectively; a 2-year influenza pandemic with an excess mortality rate of 0.5 and 1 per 1,000 in the first and second year, respectively; and, a 3-year influenza pandemic with an equal excess mortality rate of 0.5 per 1,000 in each year.

The bond is not recalibrated for scenarios (A) 3. and (B) 1. through (B) 6. as the purpose is for potential issuers to examine the effect on HE of various possible outcomes once the bond has been put in place. Given that actual outcomes will differ from those assumed in calibrating the bond, the question of interest is the extent to which the protection provided by the bond is impaired or possibly improved over a range of possible scenarios.

5 RESULTS

5.1 Base scenario

The base scenario represents the situation where actual mortality experience exactly matches the mortality assumptions used to calibrate the bond. Table 5 shows the estimated mean and variance of net claims, where net claims are defined as the aggregate claims minus the bond payoff. The mean and variance of the net claims increase when a pandemic occurs, as the higher than expected mortality causes the death of more policyholders with varying sum insured amounts. The difference between the mean of net claims for the pandemic with no bond and pandemic with bond scenarios is equal to the bond payoff, which is constant under each deterministic scenario as no sampling risk was assumed in the general population. For this reason, the estimated variance of net claims is also the same for these two scenarios.

Table 5: Estimated average and variance of life insurer's net claims.

Scenario	Estimated mean of net claims ($ Millions)	Estimated variance of net claims ($ Millions)
No pandemic with no bond	6.77	7.77
Pandemic with no bond	16.38	18.35
Pandemic with bond [a]	13.15	18.35

[a] In this scenario, the bond payoff is is equal to $3.23 million or equivalently half the principal amount because the excess mortality of 1.5 per 1,000 assumed corresponds to the middle of the range of excess mortality rate targeted in the specified hedging objectives.

The distribution of HE under the base scenario is shown in Figure 2. While the peak frequency measure is near 100%, the median HE is 115%, the mean HE is 153%, and the HE at the 5th percentile is 48%. Despite the bond being calibrated on an expected value approach, the estimated mean HE is actually 153% and not 100% as intended which indicates intrinsic over-hedging. Here the bond provides on average 53% more coverage than required. The distribution of HE shows significant skewness, with a long right tail indicating the possibility of gains. This is caused by the combination of an HE measure floored at zero, a relatively small portfolio size and a non-uniform distribution of sums insured, and is key to the effectiveness or otherwise of catastrophic mortality bonds as an alternative risk management tool. Potential bond issuers should consider using simulation to assist in the design of an appropriate hedge instrument, but need to be aware of the risk that targeting a mean HE of 100% may lead to poorer HE at the 5th percentile. The impact of portfolio size and sum insured distribution is considered further in the scenarios chosen for the sensitivity analysis.

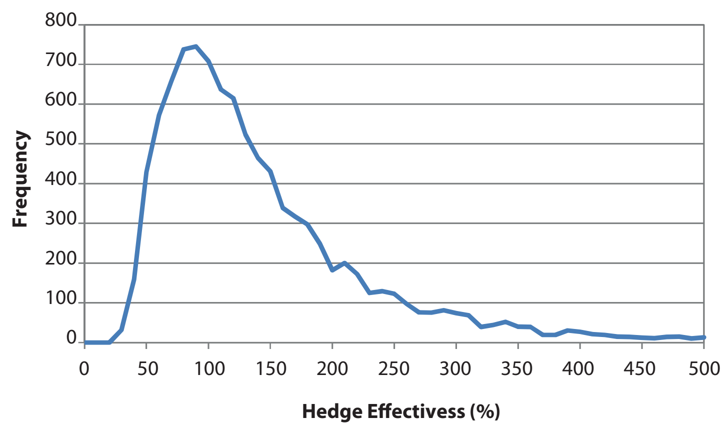

Figure 2: Estimated distribution of HE under the base scenario.

5.2 Sensitivity analysis on the characteristics of the life insurer's portfolio

In the following analysis, the HE is considered to have improved when the HE at the 5th percentile increases and the mean and median HE remain above 100%. This indicates an adequate coverage on an expected value basis, with improved protection when the outcome is significantly adverse for the insurer.

5.2.1 (A) 1. Number of policies

Table 6 shows that, as the number of policies increases, the estimated HE at the 5th percentile improves, and both the estimated median and mean HE decrease towards 100%. The latter result suggests that an expected value approach for calibrating the bond may result in a perfect mean HE for larger portfolio sizes.

Table 6: Estimated HE when varying the number of policies.

Number of policies	Estimated HE (%)		
	Mean	Median	5th percentile
20,000*	153	115	48
40,000	119	107	57
60,000	114	106	62
80,000	109	103	65
100,000	107	102	68

In this and all further tables, * indicates the base scenario.

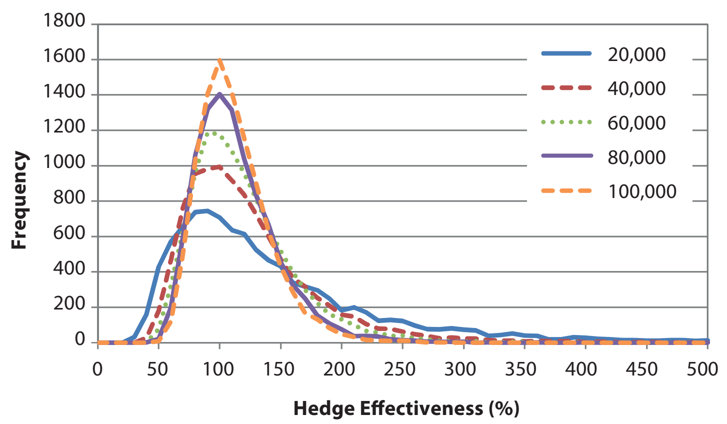

Figure 3: Estimated distribution of HE when varying the number of policies.

Figure 3 shows that the estimated distribution of HE becomes less spread, less positively skewed and more peaked as the number of policies increase. Overall, this confirms that the HE is substantially improved as portfolio size increases and clearly demonstrates the potential beneficial effect for insurers of considering pooling their portfolios when seeking protection via these instruments.

5.2.2 (A) 2. Flat sum insured distribution

Table 7 shows that a flat sum insured distribution has improved HE compared with the base scenario. The estimated mean HE is far closer to 100% while the estimated median HE also falls towards 100%. In addition, the estimated HE at the 5th percentile increases significantly.

Table 7: Estimated HE for a flat sum insured.

Distribution of sum insured	Estimated HE (%)		
	Mean	Median	5th percentile
Positively skewed*	153	115	48
Flat	114	105	63

Figure 4 shows that the estimated distribution of HE with a flat sum insured distribution has a higher peak than the base scenario. The estimated HE distribution is shown as points for the flat sum insured distribution. The HE takes only several specific values because the actual aggregate claims minus the retained claims is always a multiple of the assumed, uniform $365,000 sum insured for all policyholders, and the bond payoff is constant under each deterministic scenario.

The results suggest HE is improved when the catastrophic mortality bond is used as a hedge for a portfolio that has a flat sum insured distribution, and this demonstrates the potential for effective use of a bond in combination with a surplus reinsurance arrangement.

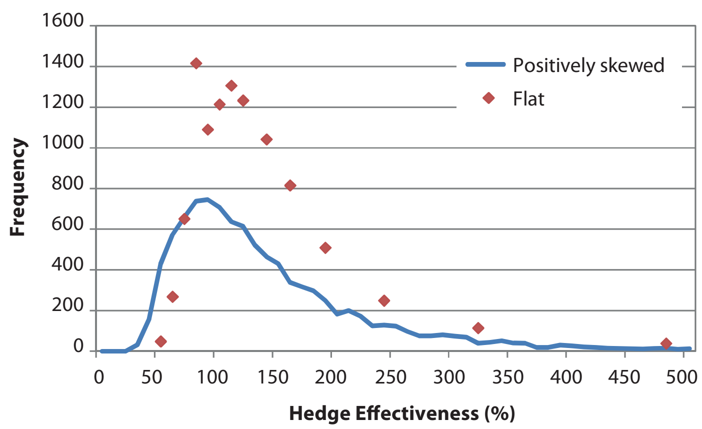

Figure 4: Estimated distribution of HE when varying the distribution of sum insured.

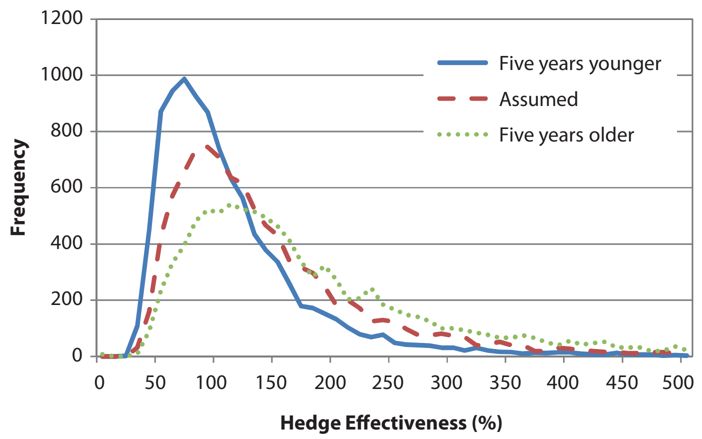

Figure 5: Estimated distribution of HE when varying portfolio composition by age.

5.2.3 (A) 3. Portfolio composition by age

Table 8 indicates that a change in the portfolio composition by age, without a change in the calibration of the bond, may significantly impact the HE. When the portfolio is 5 years younger, all the HE measures decrease in comparison to the base scenario. This is because the bond payoff remains the same while claims increase since policyholders on aggregate have increased excess mortality due to the \/\ shape affecting younger policyholders more. This represents a deterioration in HE. In comparison, all the HE measures increase with respect to the base scenario when the portfolio is five years older. Figure 5 illustrates this point.

Table 8: Estimated HE when varying the portfolio composition by age

Portfolio composition by age and gender	Estimated HE (%)		
	Mean	Median	5th percentile
Five years younger	110	88	39
Assumed*	153	115	48
Five years older	217	147	55

This finding demonstrates that the inflexibility of the bond to adjust the age (and gender) weightings for the mortality index over time may have significant consequences for the HE if the portfolio composition changes during the term of the bond. This suggests a future improvement for the structurers of such bonds could be to alter the design of bonds to better allow for this potential drift.

5.3 Sensitivity analysis of mortality assumptions

5.3.1 (B) 1. Mortality improvements

Table 9 shows that the HE is sensitive to changes in the general population mortality improvements, but not those in the insured population. As the general population mortality improvements decrease progressively in aggregate from the AGA 25-year trend to the AGA 100-year trend and then to no trend, all the examined measures of HE increase significantly as the bond payoff increases while the claims remained unchanged. The bond payoff increases as the mortality index is higher than in the base scenario. The changes

Table 9: Estimated HE when varying the general and insured population mortality improvements.

General population mortality improvement	Insured population mortality improvement	Estimated HE (%)		
		Mean	Median	5th percentile
AGA 25 year*	AGA 25 year*	153	115	48
AGA 25 year	AGA 100 year	153	115	48
AGA 25 year	None	154	114	47
AGA 100 year	AGA 25 year	174	131	54
AGA 100 year	AGA 100 year	174	131	54
AGA 100 year	None	175	130	54
None	AGA 25 year	214	161	67
None	AGA 100 year	214	161	67
None	None	215	160	66

in the insured population mortality improvements do not materially affect the HE as the chosen measure only considers the claims caused by influenza pandemic excess mortality and not baseline mortality.

Figure 6 demonstrates that the estimated distribution of HE becomes more spread as the mortality improvements weighted by the bond's age and gender calibration decrease.

5.3.2 (B) 2. Overall general population excess mortality rate

Table 10 reports the estimated HE when varying the overall general population excess mortality rate. The measures are 0% in the scenario with an excess mortality rate of 1 per 1,000 because the bond is not triggered. When the excess mortality rate increases above 2 per 1,000, the bond has already paid the entire original principal amount to the life insurer so the HE measures decrease substantially at an excess mortality rate of 2.25 per 1,000 (compared with 2 per 1,000) as the increase in claims is no longer offset by an increase in the bond payoff. These results are consistent with the intended hedging objectives of hedging against claims caused by an excess mortality rate of 1 to 2 per 1,000. As the excess mortality rate increases from 1 to 2 per 1,000, the estimated mean and median HEs decrease while the estimated HEs at the 5th percentile increase.

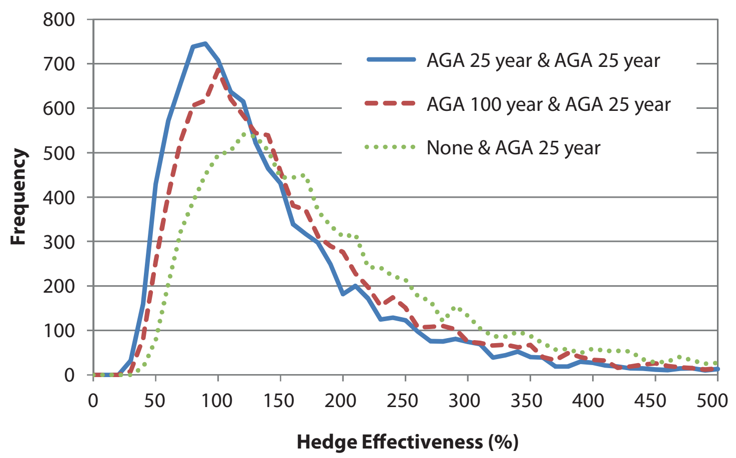

Figure 6: Estimated distribution of HE when varying the general population mortality improvements.

Table 10: Estimated HE when varying the overall general population excess mortality rate.

Overall general population excess mortality rate (per 1,000)	Estimated HE (%)		
	Mean	Median	5th percentile
1.00	0	0	0
1.25	250	125	37
1.50*	153	115	48
1.75	129	110	53
2.00	120	107	57
2.25	93	84	48

Figure 7 demonstrates that the estimated distribution of HE becomes less spread, less positively skewed and more peaked as the excess mortality rate increases from 1 to 2 per 1,000. After the excess mortality rate exceeds 2 per 1,000, the estimated distribution of HE shifts to the left and becomes more peaked. Overall, this suggests the HE is improved when the excess mortality rate is closer to the upper bound of the range used to calibrate the exhaustion point.

5.3.3 (B) 3. Age-specific distribution of excess mortality rates

Table 11 indicates that the HE is highly sensitive to changes in the age-specific distribution of excess mortality rates. As the U, W and flat shapes with the same overall general population excess mortality rate result in an increase in the mortality index less than that for the \/\ shape, the bond payoffs for these scenarios are comparatively smaller than the \/\ shape. In addition, the claims for these scenarios are also smaller, but the decrease in bond payoff is proportionally greater than the decrease in claims. This explains the lower estimated HE measures for the U, W and flat shapes in comparison to the \/\ shape. In particular, the estimated HE measures for the U shape are 0% because the bond is not triggered in this scenario as the U shape primarily affects infants and the elderly, who are given a small weighting in the mortality index.

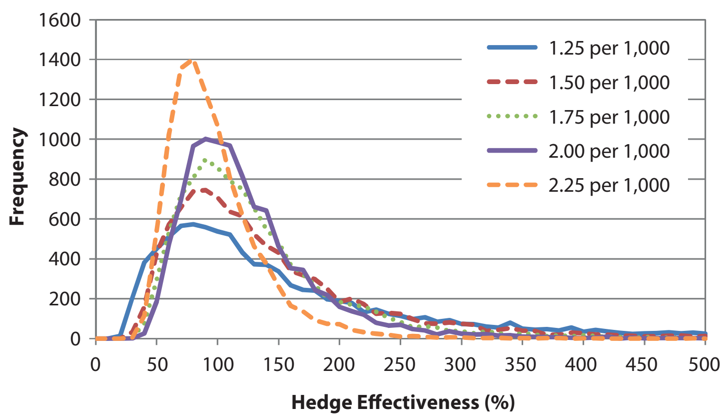

Figure 7: Estimated distribution of HE when varying overall general population excess mortality rate.

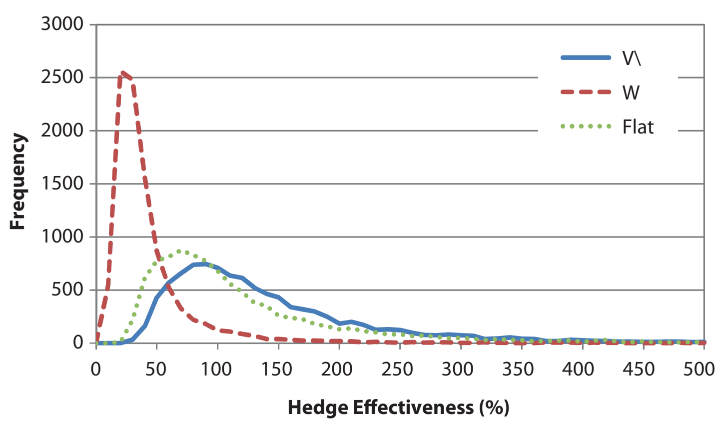

Figure 8: Distribution of HE when varying the age-specific distribution of excess mortality rates.

Table 11: Estimated HE when varying the age-specific distribution of excess mortality rates.

Age-specific distribution of excess mortality rates	Estimated HE (%)		
	Mean	Median	5th percentile
\/\ *	153	115	48
U	0	0	0
W	42	27	10
Flat	134	91	35

Figure 8 shows that the estimated distribution of HE is fairly similar between the \/\ and flat shape while it differs considerably between the \/\ and W shape. Overall, this indicates the HE of the bond is highly sensitive to the shape of age-specific distribution of excess mortality rates in a pandemic event.

5.3.4 (B) 4. Excess mortality rate ratio of insured versus general population

Table 12 demonstrates that the estimated HE measures fall as the excess mortality rate ratio increases since higher insured mortality rates result in a greater number of policyholder deaths causing claims to increase while the bond payoff remains unchanged. However the rise in the HE metrics from higher than assumed excess mortality rate ratio is small relative to the rise in those metrics when there is lower than assumed excess mortality rate ratio. Figure 9 illustrates this point.

Table 12: Estimated HE when varying excess mortality rate ratio of insured versus general population.

Excess mortality rate ratio of insured versus general population	Estimated HE (%)		
	Mean	Median	5th percentile
40%	521	250	76
60%	245	160	57
80% *	153	115	48
100%	110	89	40
120%	86	72	35

5.3.5 (B) 5. Timing of the influenza pandemic

Table 13 reports the estimated HE when varying the timing of the influenza pandemic. The HE measures for an influenza pandemic occurring in 2011 or 2012 are broadly the same. The bond payoff is made at the end of 2012 in both these scenarios since the construction of the mortality index requires two years of mortality experience. Thereafter, the estimated HE measures decrease approximately linearly every year after 2012 since the assumed positive mortality improvements decrease the general population mortality rates every year, which reduces the bond payoff. The fall in insured population mortality rates due to mortality improvements does not affect the HE as the chosen HE measure only considers the claims caused by influenza pandemic excess mortality. Figure 10 illustrates this point.

Table 13: Estimated HE when varying the timing of the influenza pandemic.

Timing of the influenza pandemic	Estimated HE (%)		
	Mean	Median	5th percentile
2011	177	132	55
2012	178	133	55
2013 *	153	115	48
2014	130	98	40
2015	107	80	33

Altogether, the findings suggest there is improved HE when the influenza pandemic occurs earlier than assumed for calibration.

5.3.6 (B) 6. Duration of the influenza pandemic and severity of waves each year

Table 14 demonstrates that the HE is sensitive to changes in the duration of the influenza pandemic and severity of waves each year, but not the order of these waves. A 2-year influenza pandemic results in lower estimated mean, median and 5th percentile HE compared with a one-year influenza pandemic. This is because a 2-year influenza pandemic has a lower bond payoff since the mortality index has decreased by one more year of mortality improvements. When the impact of influenza pandemic is spread equally over 3 years, the bond payoff is not triggered and consequently, the estimated HE measures are 0%.

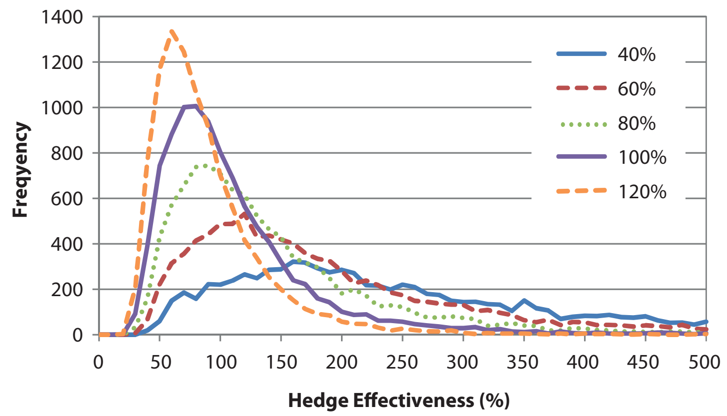

Figure 9: Estimated distribution of HE when varying excess mortality rate ratio of insured versus general population.

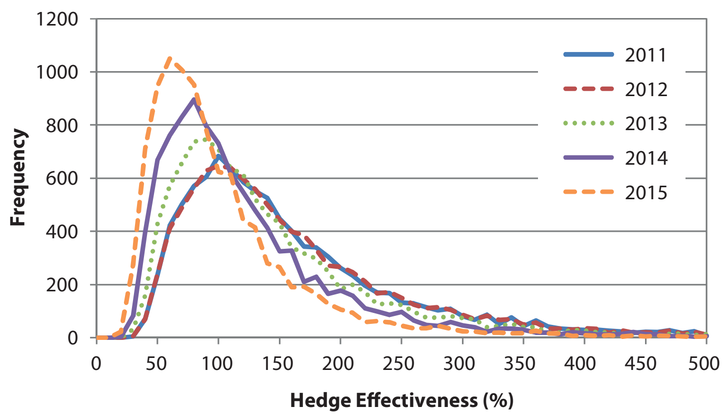

Figure 10: Estimated distribution of HE when varying the timing of influenza pandemic.

Table 14: Estimated HE when varying duration of influenza pandemic and severity of waves each year.

Duration of the influenza pandemic and severity of waves each year	Estimated HE (%)		
	Mean	Median	5th percentile
1 year *	153	115	48
2 years (stronger wave in 1st year)	128	97	40
2 years (stronger wave in 2nd year)	128	97	39
3 years (equal waves each year)	0	0	0

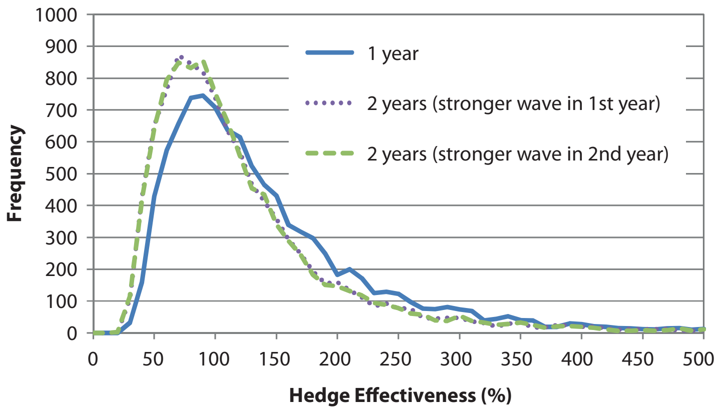

Figure 11: Estimated distribution of HE when varying the duration of the influenza pandemic and severity of waves each year.

Figure 11 indicates that the estimated distribution of HE for a one-year influenza pandemic varies somewhat from the 2-year influenza pandemic scenarios. The two year influenza pandemics with unequal waves have similar distributions. This suggests that the order of the waves does not affect HE in this case. In conclusion, it appears that the HE deteriorates for a two year influenza pandemic of equivalent total severity as the base scenario.

6 DISCUSSION AND CONCLUSION

The analysis of the base scenario finds that the hedge effectiveness of catastrophic mortality bonds is highly variable. Although the bond is intended to provide strong hedge effectiveness for that scenario the actual estimated mean hedge effectiveness implies substantial over-hedging. This is primarily attributed to the positively skewed distribution of sum insured and the relatively small portfolio size assumed in the base scenario.

The sensitivity analysis on the characteristics of the life insurer's portfolio suggests that catastrophic mortality bonds have improved hedge effectiveness under certain circumstances. Catastrophic mortality bonds appear to be viable alternative risk management tools for large portfolio sizes particularly where the distribution of sum insured is more uniform, and when the life insurer's underlying exposure remains relatively stable.

As the number of policies in the portfolio increases, the overall hedge effectiveness improves significantly. In addition, the estimated mean hedge effectiveness converges to 100% as intended when using the expected value approach for calibrating the bond. This is coherent with the current use of catastrophic mortality bonds as a risk management tool for large, globally diversified insurers and reinsurers. In comparison, reinsurance or retrocession coverage is likely to remain a significantly better risk management tool compared to catastrophic mortality bond for smaller insurers or reinsurers since there is significant basis risk and variation in the hedge effectiveness at smaller portfolio sizes. However, it may be possible for smaller insurers or reinsurers to pool their exposures and issue a catastrophic mortality bond for the aggregated portfolio.

Although age, gender, country and financial exposure in the form of sum insured can be calibrated for catastrophic mortality bonds to match the life insurer's underlying exposure, the hedge effectiveness is highly sensitive to changes in the portfolio composition. As the age and gender weightings of the bond are fixed at issuance, any change in the composition of the life insurer's portfolio could have a substantial impact on the hedge effectiveness. In particular, a change in the age composition is likely to have a greater impact than a change in gender composition because historical influenza pandemic mortality experience suggests excess mortality rates vary considerably across age, but not by gender. Consequently, factors that may affect the portfolio composition, such as a change in marketing and advertising strategy, will have serious ramifications for the hedge effectiveness of catastrophic mortality bonds.

Catastrophic mortality bonds provide lower variation in hedge effectiveness for portfolios where the distribution of sum insured is more uniform. A possible hedging strategy stemming from this result is the combined implementation of surplus reinsurance and catastrophic mortality bonds. Firstly, surplus reinsurance should be used to transfer the volatility in the sum insured exposure above a specified retention. Secondly, a catastrophic mortality bond should be used to cover the life insurer's retention. The surplus reinsurance effectively reduces the spread of the distribution of sum insured for which the catastrophic mortality bond is used to provide coverage.

The sensitivity analysis on the mortality assumptions highlights the significant uncertainty surrounding the basis risk and hedge effectiveness of catastrophic mortality bonds. This is due to the inherent uncertainty regarding future mortality rates, particularly in a pandemic scenario where the actual epidemiological characteristics are impossible to predict. In general, catastrophic mortality bonds provide improved hedge effectiveness when: the general population mortality improvements are lower than assumed; the overall general population excess mortality rate is at the upper bound of the range used to calibrate the exhaustion point; the age-specific distribution of excess mortality rates follows the assumed shape; the excess mortality rate ratio of insured versus general population is lower than assumed; the influenza pandemic occurs at the start of the risk period; and, the duration of the influenza pandemic is 1 or 2 years.

The research could be extended to analyse the hedge effectiveness for aggregate mortality exposure across a range of life insurance products. For example, group life portfolios, traditional products, retirement income products and other risk products could be examined. Other methods of calibrating the characteristics of catastrophic mortality bonds could also be explored as there seems to be minimal existing literature regarding this area. A detailed investigation into the calibration of catastrophic mortality bonds should provide a better understanding of how to optimise the hedge effectiveness.

Furthermore, a holistic approach to enterprise risk management may consider the interaction of catastrophic mortality bonds with existing risk management strategies such as reinsurance. For example, the previously suggested hedging strategy of pooling portfolios for a bond issuance and/or using both surplus reinsurance and catastrophic mortality bonds to hedge against catastrophic mortality events could be explicitly investigated.

APPENDIX

Figure 12: Basic catastrophic mortality bond transaction structure.

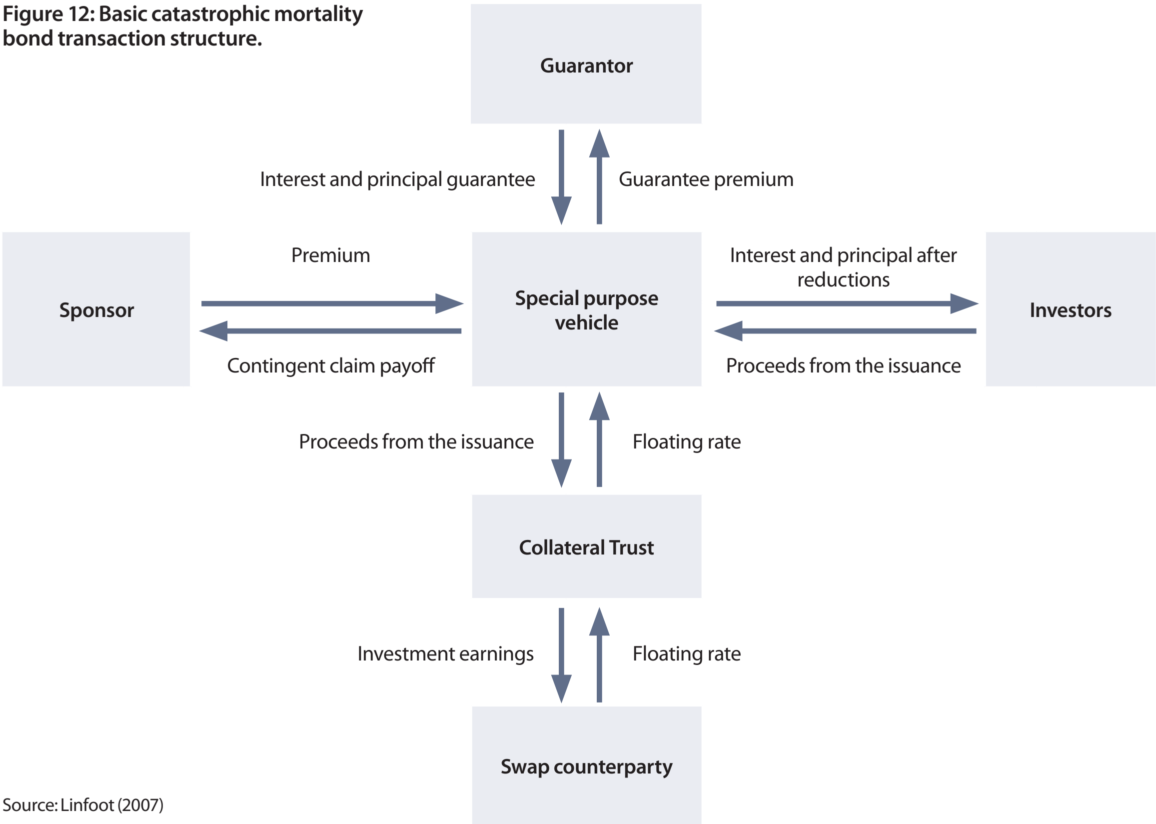

Source: Linfoot (2007)

Additional technical notes:

The general and insured population excess mortality rate for five year age group i, $GPEMR_i$ and $IPEMR_i$, that are applied over the assumed duration of the influenza pandemic are calculated as follows:

$$GPEMR_i = OGPEMR * R_i$$
$$IPEMR_i = GPEMR_i * EMRR$$

Where:

$GPEMR_i$	=	The general population excess mortality rate for 5-year age group *i*;
$OGPEMR$	=	The overall general population excess mortality rate
R_i	=	The ratio of general population excess mortality rate for five year age group *i* to the overall general population excess mortality rate
$IPEMR_i$	=	The insured population excess mortality rate for 5-year age group *i*
$EMRR$	=	The excess mortality rate ratio of insured versus general population.

The catastrophic mortality bond payoff at the end of measurement period in calendar year *t*, $CMBP_t$, can be expressed as followed:

$$CMBP_t = P \times R_t^C$$

$$R_t^C = Max\left(\frac{Index_t^C - A^C}{E^C - A^C} - R_{t-1}^C, 0\right)$$

Where:

$CMBP_t$	=	The catastrophic mortality bond payoff at the end of measurement period in calendar year *t*
P	=	The original principal amount
R_t^C	=	The principal reduction factor for measurement period ending in calendar year *t* and country *C* where $R_t^C = 0$ for the start of the risk period and $0\% \leq \sum_t R_t^C \leq 100\%$
$Index_t^C$	=	The mortality index for measurement period ending in calendar year *t* and country *C*
A^C	=	The attachment point for country *C*
E^C	=	The exhaustion point for country *C*.

The mortality index for measurement period ending in calendar year *t* and country *C*, $Index_t^C$, can be expressed as:

$$Index_t^C = \frac{\bar{q}_t^C}{\bar{q}_{\text{reference years}}^C}$$

$$\bar{q}_t^C = \frac{1}{2} * (q_t^C + q_{t-1}^C)$$

$$q_t^C = \sum_x (w_{x,m}^C * q_{m,x,t}^C + w_{x,f}^C * q_{f,x,t}^C)$$

Where:

$Index_t^C$	=	The mortality index for measurement period ending in calendar year *t* and country *C*
$\bar{q}_t^C$	=	The two year average mortality rate over measurement period ending in calendar year *t* for country *C* where the reference years correspond to the two years before the start of the risk period
q_t^C	=	The annual mortality rate for country *C* in calendar year *t*
$q_{m,x,t}^C$	=	The mortality rate for males of country *C* and age group *x* in calendar year *t*
$q_{f,x,t}^C$	=	The mortality rate for females of country *C* and age group *x* in calendar year *t*
$w_{x,m}^C$	=	The weight applied to male mortality rates of country *C* and age group *x*
$w_{x,f}^C$	=	The weight applied to female mortality rates of country *C* and age group *x*.

The future average sum insured in year *t*, ASI_t, can be expressed by the following formula:

$$ASI_t = ASI_s * \prod_{x=s}^{t-1} (1 + TUR * AIR_x)$$

Where:

ASI_t	=	The future average sum insured in year t
ASI_s	=	The past average sum insured in year s, where $s < t$
TUR	=	The take up rate
AIR_x	=	The annual inflation rate in year x.

Bibliography

Australian Bureau of Statistics (ABS). (2010). Life Tables, Australia, 2007-2009 *Cat. No. 3302.0.55.001*. Canberra: Australian Bureau of Statistics.

Australian Government Actuary (AGA). (2009). Australian Life Tables 2005-07. Canberra: Australian Government Actuary.

Australian Prudential Regulation Authority (APRA). (2007). APRA's Pandemic Stress Test of the Insurance Industry *APRA Insight* (Vol. Issue 3, 2007). Sydney: Australian Prudential Regulation Authority.

Bagus, G. (2007). Taming the Cat: Securitizing pandemic losses offers life insurers a guaranteed way to fund potential claims. *Best's Review,* September 2007, 92-94.

Baumgart, C., Lempertseder, R., Riswadkar, A., Woolnough, K., & Zweimüller, M. (2007). Influenza Pandemics. *CRO briefing Emerging Risks Initiative – Position Paper*: CRO Forum.

Cairns, A. J. G., Dowd, K., Blake, D., & Coughlan, G. D. (2011). Longevity Hedge Effectiveness: A Decomposition. *Discussion Paper*: Pensions Institute.

Chan, P., Gale, A. P., Mak, A. T., Muir, P. R., Murphy, C. D., O'Carra, S., . . . Wildgoose, J. Y. (2001). Institute of Actuaries of Australia Mortality Investigation IA 95-97 M and F: Graduated Mortality Tables: Institute of Actuaries of Australia.

Coughlan, G. D., Allah, M. K., Ye, Y., Kumar, S., Cairns, A. J. G., Blake, D., & Dowd, K. (2011). Longevity hedging 101: A framework for longevity basis risk analysis and hedge effectiveness. *North American Actuarial Journal, 15*(2), 150–176.

Cowley, A., & Cummins, J. D. (2005). Securitization of life insurance assets and liabilities. *Journal of Risk and Insurance,* 72(2), 193-226.

Cox, S. H., & Hu, Y. (2004). Modeling mortality risk from exposure to a potential future extreme event and its impact on life insurance. Atlanta: Department of Risk Management and Insurance J. Mack Robinson College of Business.

Cox, S. H., & Lin, Y. (2007). Natural Hedging of Life and Annuity Mortality Risks. *North American Actuarial Journal,* 11(3), 1–15.

Craig, J., & Dublin, L. (1919). The Influenza Epidemic of 1918 *Transactions* (Vol. 20, pp. 134-156): Actuarial Society of America.

Cummins, J. D., Lalonde, D., & Phillips, R. D. (2004). The Basis Risk of Catastrophic-Loss Index Securities. *Journal of Financial Economics,* 71(1), 77–111. doi: 10.1016/s0304–405x(03)00172–7.

Dauer, C. C., & Serfling, R. E. (1961). Mortality from influenza, 1957–1958 and 1959-1960. *American Review of Respiratory Disease,* 83, 15-28.

Dreyer, A., Kritzinger, G., & Decker, J. D. (2007). Assessing the Impact of a Pandemic on the Life Insurance Industry in South Africa. Paper presented at the 1st International Actuarial Association Life Colloquium, Stockholm.

Faulds, T. G., & Bridel, W. (2009). Considerations for the Development of a Pandemic Scenario Research Paper. Ottawa: Canadian Institute of Actuaries.

Glezen, W. P. (1996). Emerging Infections: Pandemic Influenza. *Epidemiologic Reviews,* 18(1), 64–76.

Harrington, S., & Niehaus, G. (1999). Basis Risk with PCS Catastrophe Insurance Derivative Contracts. *The Journal of Risk and Insurance,* 66(1), 49–82.

Hayes, G. K. (2008). Stochastic Solvency Testing in Life Insurance. PhD thesis Australian National University, Canberra.

Huynh, A., Bruhn, A., & Browne, B. (2013). A Review of Catastrophic Risks for Life Insurers. *Risk Management and Insurance Review,* 16(2), 233–266. doi: 10.1111/rmir.12011.

Li, J. S. H., & Hardy, M. R. (2011). Measuring Basis Risk in Longevity Hedges *North American Actuarial Journal,* 15(2), 177–200.

Linfoot, A. (2007). Financing Catastrophic Risk: Mortality Bond Case Study: Scottish Re.

Little, J. (1919). Discussion of the paper 'The influenza epidemic of 1918, Craig J and Dublin L' *Transactions* (Vol. 20, pp. 493–496): Actuarial Society of America.

Mead, F. B. (1919). Discussion of the paper 'The influenza epidemic of 1918, Craig J and Dublin L' *Transactions* (Vol. 20, pp. 496–498): Actuarial Society of America.

Metropolitan Life Insurance Company. (1976). Mortality from influenza and pneumonia. *Metropolitan Life Insurance Company Statistical Bulletin,* 57, 2–7.

Murray, C. J. L., Lopez, A. D., Chin, B., Feehan, D., & Hill, K. H. (2006). Estimation of potential global pandemic influenza mortality on the basis of vital registry data from the 1918–20 pandemic: a quantitative analysis. *The Lancet,* 368(9554), 2211–2218. doi: 10.1016/s0140-6736(06)69895–4.

Ngai, A., & Sherris, M. (2011). Longevity risk management for life and variable annuities: The effectiveness of static hedging using longevity bonds and derivatives. *Insurance: Mathematics and Economics,* 49(1), 100–114. doi: 10.1016/j.insmatheco.2011.02.009.

Nguyen-Van-Tam, J. S., & Hampson, A. W. (2003). The epidemiology and clinical impact of pandemic influenza. *Vaccine,* 21(16), 1762–1768. doi: 10.1016/s0264-410x(03)00069-0.

Palmer, N. (2009). Life Insurance: Using Experience Studies to Strengthen and Grow the Industry.

Potter, C. W. (2001). A history of influenza. *Journal of Applied Microbiology,* 91(4), 572–579. doi: 10.1046/j.1365-2672.2001.01492.x.

Richards, S., & Jones, G. (2004). *Financial aspects of longevity risk*. London: Staple Inn Actuarial Society.

Simonsen, L., Clarke, M. J., Schonberger, L. B., Arden, N. H., Cox, N. J., & Fukuda, K. (1998). Pandemic versus epidemic influenza mortality: A pattern of changing age distribution. *Journal of Infectious Diseases,* 178, 53–60.

Simonsen, L., Clarke, M. J., Williamson, G. D., Stroup, D. F., Arden, N. H., & Schonberger, L. B. (1997). The impact of influenza epidemics on mortality: introducing a severity index. *Am J Public Health,* 87(12), 1944–1950. doi: 10.2105/ajph.87.12.1944.

Stracke, A. (2007). Influenza Pandemic – The Impact on an Insured Lives Life Insurance Portfolio. *Risk Insights,* 11(4), 18–21.

Swiss Re Capital Markets. (2008). Supplement Listing Particulars No. 1 to Offering Circular Dated January 18, 2008 Vita Capital III Ltd. : Swiss Re Capital Markets.

Toole, J. (2007). *Potential Impact of Pandemic Influenza on the U.S. Life Insurance Industry*: Society of Actuaries.

United Nations. (1999). The World at Six Billion Retrieved 17 August, 2011, from http://www.un.org/esa/population/publications/sixbillion/sixbilpart1.pdf.

US Census Bureau. (2000). Historical National Population Estimates: July 1, 1900 to July 1, 1999 Retrieved 10 August, 2011, from http://www.census.gov/popest/archives/1990s/popclockest.txt.

US Census Bureau. (2011a). Preliminary Annual Estimates of the Resident Population for the United States, Regions, States, and Puerto Rico : April 1, 2000 to July 1, 2010 Retrieved 12 August, 2011, from http://www.census.gov/popest/eval-estimates/eval-est2010.html.

US Census Bureau. (2011b). Total Midyear Population for the World: 1950-2050. *International Data Base* Retrieved 18 August, 2011, from http://www.census.gov/population/international/data/idb/worldpoptotal.php.

US Department of Health and Human Services. (2011). Pandemics and Pandemic Threats since 1900 Retrieved 16 August, 2011, from http://www.pandemicflu.gov/general/historicaloverview.html.

Viboud, C., Miller, M., Olson, D., Osterholm, M., & Simonsen, L. (2010). Preliminary Estimates of Mortality and Years of Life Lost Associated with the 2009 A/H1N1 Pandemic in the US and Comparison with Past Influenza Seasons. *Public Library of Science Currents, 2*. doi: 10.1371/currents.RRN1153.

Weisbart, S. (2006). Pandemic: C*an the Life Insurance Industry survive the Avian Flu?* New York: Insurance Information Institute.

Woolnough, K., Ivanovic, B., Kramer, S., & Busenhart, J. (2007). Pandemic Influenza: A 21st century model for mortality shocks. In T. Dickenson (Ed.). Zurich: Swiss Re.

World Health Organisation. (2005). *Avian influenza: assessing the pandemic threat*. Geneva: World Health Organisation.

Zeng, L. (2000). On the Basis Risk of Industry Loss Warranties. *Journal of Risk Finance,* 1(4), 27–33.

Alex Huynh
alex.huynh@booz.com
Alex Huynh is a Senior Consultant at Booz and Company working in the Sydney office. Previously, he completed a Bachelor of Actuarial Studies (Honours) at the Australian National University.

Bridget Browne
bridget.browne@anu.edu.au
Bridget Browne is a Senior Lecturer in Actuarial Studies at the Australian National University in Canberra. Previously she was Global Life Chief Pricing Actuary at Partner Reinsurance.

Aaron Bruhn
aaron.bruhn@anu.edu.au
Aaron Bruhn lectures in Actuarial Studies at the Australian National University. Previously he worked in Life Insurance and as a Principal Adviser for the New Zealand Ministry of Defence.

A pilot survey of actuarial graduates' views on their education

DR ADAM BUTT, JOHN EVANS, JIM FARMER AND DR DAVID PITT

Dr Adam Butt

John Evans

Dr David Pitt

ABSTRACT

This paper presents the results of a pilot survey of recent graduates of some Australian university actuarial programs. The survey aimed to shed light on graduates' views relating to their education since leaving university. The survey that we report on here has been designed as a pilot to inform the development of a more substantial survey to be conducted in future. It does, however, still provide some interesting insights. The findings from our work can be used by those currently reviewing the actuarial education programs in Australia. The broad results indicate that around 50% of students found their Part I courses covered technical material useful in their employment, and well over 50% of graduates considered non-technical training as very important for their readiness for employment. The Part II results indicated a much higher satisfaction with the content than the Part I courses in preparing students for employment. The Part III courses were not rated as highly as the Part I and II courses and the teaching quality was particularly criticised. Overall, the results indicated that the graduates surveyed found their Part I and II courses valuable but wanted greater training in the non-technical skills needed in employment, particularly communication skills, and while they found the Part III courses applicable to their employment, they expressed dissatisfaction with the education process.

KEYWORDS

actuarial education, graduate survey

1. INTRODUCTION

This paper presents the results from a pilot survey conducted on a group of selected graduates from the past five years of actuarial programs at the Australian National University, Macquarie University, the University of Melbourne and the University of New South Wales. The purpose of the survey was twofold: firstly to ascertain the views of graduates on aspects of their university and post-university education and secondly to test the questions asked before engaging in a survey of a larger group.

Syllabus development and renewal are an important part of the work of a professional body. Currently the Actuaries Institute is conducting a holistic review of education programs for actuarial students in Australia. This review is being informed in part by the Actuaries Institute's Capability Framework which outlines a number of generic skills or competencies deemed essential for the practicing actuary. These skills include written and verbal communication, problem solving, ethics, teamwork and leadership and computer programming abilities. At the same time, university executives are also working to implement the recommendations in the recently developed Australian Qualifications Framework (AQF). This framework provides an outline of the many requirements that are regarded as essential for Australian degrees at the bachelors, masters and PhD levels and includes research, technical and generic competencies as part of its recommendations. These activities are a primary reason for publishing the results of the pilot survey before running the larger survey; we feel that even the limited scope of the results of the pilot survey will be beneficial to those undertaking these activities. We also hope that the publication of the results of the pilot survey might encourage actuarial groups in other jurisdictions to undertake similar work.

Past research into graduate views on actuarial education is very limited, and is hence a motivating aspect of this research. Chu et al. (2010) report on a survey of actuarial graduates from a single Australian university. The questions asked in the survey in Chu et al. (2010) are broader in their assessment of the usefulness of actuarial education and do not focus on specific syllabus items like the survey described in this paper. Butt et al. (2011) report on a series of interviews with recent graduates from a single Australian university. The interview questions asked in Butt et al. (2011) can be thought of as a precursor to the survey described in this paper. Respondents in both Chu et al. (2010) and Butt et al. (2011) indicate the importance of an education that is broad and contextualised in problems students will face upon graduation, with the respondents in Butt et al. (2011) expressing a concern that actuarial education may to some degree sacrifice the contextual elements of education due to the abundance of technical content. These concerns are consistent with surveys of employers of actuaries, who express concerns at the weaknesses of actuaries in communication, interpersonal relationships and strategic/commercial thinking (see Ipsos 2010; Howes 2011).

After this introduction, Section 2 provides an outline of the literature in the area of surveys on the views that students and graduates have of their university education as relevant to our research. Section 3 describes the survey used in this paper. Section 4 summarises the results of our survey and Section 5 concludes the paper.

2. A REVIEW OF RELATED LITERATURE

In this section we report on a review of the relevant literature. Before doing this it is relevant to mention a well-known national survey, the Australian Graduate Survey (AGS) which incorporates the Course Experience Questionnaire (CEQ). The AGS is 'the national census of newly qualified higher education graduates. Conducted annually since 1972, the AGS surveys new graduates from all Australian universities, and a number of higher education institutes and colleges, approximately four months after they complete the requirements for their awards.' The Course Experience Questionnaire (CEQ) probes key elements of the higher education experience relevant to coursework graduates, focusing largely on their perceptions of course quality, their self-rated skill levels, and their overall satisfaction with their course.' The main limitation of using the CEQ for informing curriculum development in a particular discipline is that it is a broad survey conducted nationwide and as a result the data are not easily subdivided to the level required for analysis at the level of a specialised degree program like actuarial studies. It also suffers from low response rates. Given these limitations, we report on some different more specialised research output that uses surveys of students and graduates.

Relevant literature includes studies directed at different groups of people. There are surveys of current students in university programs, students who have just graduated from a university program and former students who have graduated some time ago at the date of survey. The nature of the questions asked in such surveys also differs: they can be classified as questions designed to assess the skills and competencies developed in the education program with a view to preparing students for working in a range of fields as opposed to questions relating to how the education program prepared the student for entry into a specific profession. Finally, surveys can focus on the technical content of an education program and/or the ability of the program to develop more generic skills.

We find that published research into the perceptions of professional exams is scarce relative to research into the content of university degree programs which may help students acquire knowledge needed for professional qualifications. We also find that research on the perceptions of relatively recent graduates is very rare while research into the views of current and graduating students is more common. Research on broad and overall perceptions of actuarial education is limited to the brief discussion made in the introduction to this paper.

Given these differences in the amount of available research literature by survey design and intended audience and our attempt to provide background to our study which focusses on graduates' views of professional education, we provide a description of four papers chosen to inform our work in order to provide some background to the discussion that follows in Sections 3 and 4. The first three papers are from the accounting field with the fourth from a science field. The bias in our review towards accounting research is partly due to the large amount of material available in this area, but also because papers from that field most closely match what we are trying to achieve in our study. The papers are relatively recent and have a local flavour, but those wishing to read further will find all four papers discussed below contain extensive pointers to a large body of prior research.

De Lange, Jackling and Gut (2006) provide one of the rare surveys of graduates, rather than current students or just completing students. The study surveyed students graduating from accounting programs in two Victorian universities in 2001 to 2003, described as the 'previous three years', placing the survey date in 2004. Respondents 'perceived communication- and analytical-based skills as being the most important qualities required for a successful accounting career.' Respondents were also asked for their perceptions of how well particular skills were covered in their degree and how much coverage there should have been. For every single skill surveyed, which included both generic and technical accounting skills, the average scores indicated the skill should have received more coverage than it did. This is an interesting reversal on the common wisdom that current students usually complain that there is too much material in their degree. The greatest discrepancies between actual coverage and desired coverage were observed for interpersonal skills, oral expression and computing skills.

Keneley and Jackling (2011) report on a survey of second year accounting students at an Australian university seeking their perceptions of the generic skills they acquired in their university accounting studies. While the study found that students perceived that generic skills were developed, international students were more likely than domestic students to rate the unit as developing some generic skills. The study notes the possible confounding factor that international students may have had less exposure to these skills in prior studies. A student is more like to find the unit develops a particular skill if they have not had prior training in that skill. The differences tended to occur on 'behavioural skills', such as questioning of accepted wisdom, thinking creatively and thinking and reasoning logically, rather than on 'cognitive generic skills' such as problem-solving and adapting knowledge to new situations.

Kavanagh and Drennan (2008) report on a survey of '322 graduating accounting students in three universities in Australia and 28 practitioners across a number of organisations and industries who employ accounting graduates'. They find that the graduating students do recognise the importance of a range of generic skills. However, students reported that their degrees fall short on delivering some career skills that they perceived as important, particularly self-motivation, professional attitude, oral communication, decision-making and continuous learning. They also find that while graduating students and employers do tend to regard the same types of skills as important, there are some significant differences between the two groups in the relative importance of some factors. Students placed more emphasis on technical skills, decision-making, critical thinking and self-motivation. Employers placed more emphasis on general business awareness, knowledge of ethics and the profession, an ability to work across disciplines and interpersonal skills. Both groups rank oral communication skills as important and this was one skill that students identified as not being well developed by their degrees. While the study did not explicitly ask respondents to think about whether particular skills were best developed during the degree or in the workplace, there was a strong perception that the employers surveyed tended to the perhaps optimistic view that graduates should arrive for work fully skilled.

A general feature of research in the accounting field is a perception that many of the 'traditional' technical skills required of accountants can now be reliably performed by software, so that accounting degrees should be able to spend less time on routine technical matters and more time on higher order thinking and generic skills. There is disagreement as to whether this has been achieved yet – the disagreement possibly arising from differences in the extent to which different programs have implemented changes in this area. While less commonly encountered, there is also a view that developments in software create extra job opportunities in fraud detection, or at least that software creates different opportunities for fraud so that different skills sets are required to detect it.

Matthews and Hodgson (2012) report on trials of a purpose-built survey administered to graduating biomedical students at University of Queensland and Monash University. The study was carried out against the background of devising tools for assessing courses against the Australian Quality Framework. The survey found students recognised the importance of writing, communication and teamwork skills. The respondents felt these skills had been addressed in their course, felt their skill levels had increased and felt confident in applying these skills. Results were broadly consistent across the two institutions. They also note that one institution had implemented an initiative to improve quantitative skills, the result being that students from that institution were more aware of the importance of those skills and were more likely to agree these skills were addressed, but continued to report low confidence with these skills comparable to the other institution. The authors note:

> Student surveys can provide quick, low-cost and meaningful information about student acquisition of knowledge and skills which can be used to inform curriculum development. Surveys of students' experiences and learning in the Higher Education sector have become important sources of data for governments seeking measures of institutional success and political accountability.

The following is a brief summary of issues that recur over several of these papers, or in some cases, issues that are surprising by their absence.

1. Graduate surveys often struggle to obtain a respectable response rate, giving the risk that they summarise the views of an unrepresentative sample of the profession involved. In general terms, while a survey of current students might get a response rate of about 1 in 2, a survey of graduates is lucky to get 1 in 4.
2. Results are diffuse. Low response rates tend to result in studies claiming the data suggest particular ideas rather than being able to state definite conclusions. Findings vary by profession and even between papers within a single profession.
3. Likert scales predominate. The 5-point Likert scale is by far the most common, but there is occasional use of a 4-point scale which attempts to force respondents to take a side rather than selecting a neutral response. One thing that distinguishes our survey from the general pattern is the frequent provision of text-boxes allowing respondents to provide further information.
4. Surveys inevitably provide a list of skills and ask respondents questions about them. It could be argued this amounts to 'leading the witness' and that instead respondents should be asked an open-ended question such as 'tell us which generic skills you developed.' While we have not found any papers explicitly addressing this issue, we feel the latter approach is problematic, since asking graduates about what they learned in their studies or what skills should be in the syllabus inevitably results in responses about discipline-specific knowledge and skills. The only generic skill likely to be mentioned is use of spreadsheets. To obtain information about generic skills, we asked respondents about the importance of, and education of, an explicit list of generic skills inferred from the Actuaries Institute's Capability Framework and also asked for open-ended comments.
5. A common thread is the tension between (some) employers who feel that graduates should come with mastery of all the generic skills required in the workplace, so that the employer need not spend any time on further training, and (some) academics who feel there is not enough solid theory in how these skills can be developed to warrant their inclusion in a university course. Perhaps also both parties take the view that helping people develop these skills can be very time consuming, so both are anxious to avoid that cost.
6. Studies are predominantly carried out by academics in the professions being surveyed, and this may lead to some biases in the approaches taken and arguments presented. For example, many papers mention that a particular program's graduates are having greater difficulty finding employment. To a disinterested spectator, perhaps the first obvious questions are whether the program is admitting too many students and thus producing too many graduates, or whether there are too many universities offering this program. However, the research found does not ask this question, but instead asks: Can we adjust the skills our students learn to make them more employable? Some papers suggest improving graduates' generic skills is one way to do this, since employers increasingly seem to be seeking work-ready graduates. However, perhaps employers' increasing tendency to seek work-ready graduates is merely the natural response of being in a buyers' market. With the

number of graduates exceeding the number of jobs, employers can be more selective. Perhaps placing greater emphasis on generic skills will alter which graduates find relevant work, but not actually increase the number finding relevant work. Anecdotal evidence suggests that there is a significant oversupply of actuarial graduates who wish to commence an 'actuarial' job in Australia compared with the number of positions available.

7. The actuarial profession has a long history of robust discussion about the content of the syllabus of the professional exams with discussion spanning all levels of the qualification process and surveying student experiences. To actuaries, this seems natural. When reading papers from other professions, it was odd that some seemed a little self-conscious about surveying students, graduates or employers on what skills a university program should concentrate on, and felt a need to justify doing so. Justification often referred to the need to consult all stakeholders.
8. While members of our profession frequently argue and complain about exactly what should be in the syllabus of the professional exams, we should perhaps count ourselves as lucky to have a single agreed national syllabus that has been relatively stable over time. By contrast, in other professions, a profession-based survey of graduates' abilities may be confounded by major differences in knowledge sets between graduates from different institutions.

3. SURVEY DESCRIPTION

Before any potential survey respondents were contacted, ethics approval was obtained from the relevant committees of all authors' institutions.

A. Sample selection

A primary aim of the pilot survey was to test out the survey structure and questions, to determine whether any questions were misinterpreted and to see whether responses to open-ended questions suggested areas where further explicit questions might prove useful. The outcome of this aim is discussed at the start of Section 4. Given this purpose, we ensured that the survey was undertaken by a range of respondents, but did not explicitly ensure that the respondents were a representative sample. Hence caution is called for if attempting to extrapolate from the results of this pilot survey. Some key demographic details of the respondents are provided in Section 4.

The sample group was selected by each initial investigator compiling a list of recent graduates known to them. Care was taken to include both domestic students and international students, including some whose first language was not English. One possible source of bias is the tendency for the most well-known graduates to be the students who were more successful in their studies, so 'struggling students' may be under-represented. Graduates must have completed an actuarial degree to be included. The sample did not, for example, include graduates who commenced a double degree or combined degree that included actuarial studies, but who discontinued the actuarial half of that degree. The sampling approach could be considered to be a convenience approach, albeit with some clustering to improve the representativeness of the sample.

In total 55 graduates were invited to participate in the survey.

B. Survey implementation

The survey was implemented via the SurveyMonkey website. The investigators each emailed potential respondents inviting them to take part and providing the URL to the survey. Online surveys tend to produce lower error rates than paper-based surveys, since they can build in the logic to flag errors as they occur, allowing the respondent to fix them before submitting the survey. The survey incorporates branch logic, so that questions may be hidden if answers to earlier questions indicate they are not relevant.

The invitation to respondents was made in late April and early May of 2013. One reminder was emailed to potential participants before the survey was closed to additional responses at the end of May 2013.

C. Survey questions

For reasons of brevity, the full list of questions asked in the survey is not presented here, although interested parties are welcome to contact the corresponding author who will gladly provide them. This sub-section provides a brief summary of the content of the survey. It should also be noted that the figure and table titles in Section 4 represent the exact wording of the questions asked to the survey respondents.

The survey commences with a range of demographic questions which may help detect any strong differences between distinct sub-groups. This includes such information as age, field and country of employment, length of employment, membership of actuarial bodies and detailed data concerning progress through the Australian professional exams.

It then moves on to the major issues we are investigating. For each of Parts I, II and III of the professional exams that respondents have attempted, they are asked to rate the following factors.

- perceived usefulness of material in their current employment
- strengths or weaknesses of factors such as syllabus and teaching quality
- the size of the syllabus.

The survey then moves onto items mentioned in the Institute's Capability Framework, seeking information about written and oral communication skills, leadership, problem-solving, professionalism, ethics, strategy and teamwork. For each of these areas, respondents are asked for their perceptions on the importance of these skills, how well these skills were developed over Part I, II and III and elsewhere, and where respondents believe these skills should best be developed.

Throughout these questions, there are numerous options allowing respondents to include additional open-ended responses that clarify the reasons for their ratings. These typically took the form of 'Please comment on your answer to this question'.

Since this was a pilot survey, it finished with an open-ended question asking respondents to identify any issues they had with the survey design. The results of the survey are discussed in the next section.

4. SURVEY RESULTS

Of the 55 graduates who were invited to participate in the pilot survey, 34 commenced the survey and 28 completed the survey. Of these 34 respondents:

- 5 had attempted Part I only, while 15 had attempted at least one Part III module
- 10 were involved in an actuarial education module at the time of the survey, whilst the other respondents had completed the last of their formal actuarial education over the last five years or so
- 21 were employed in Australia, with the majority of the remainder in Asia
- a wide variety of employment areas were represented, with the most common being 8 in general insurance and 7 in life insurance.

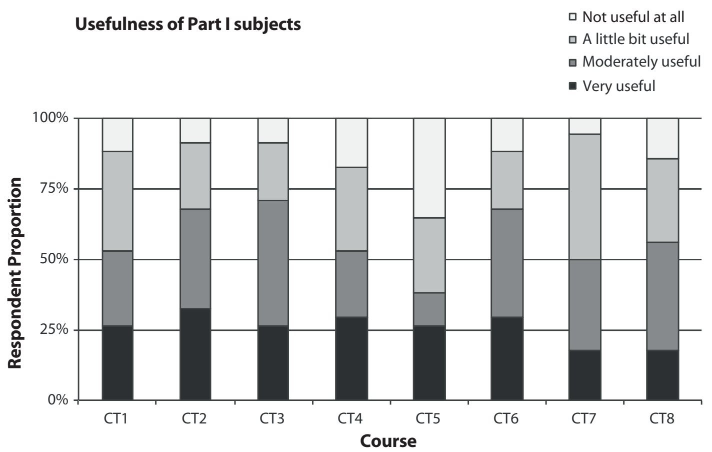

Figure 1: Rate the usefulness of the Part I subjects to your current employment role.

More detailed demographic information on the respondents can be found in the Appendix.

The figures and tables in this section outline the most interesting results of the pilot survey. In addition, where appropriate representative quotes from open ended questions are presented in italics. Due to the relatively small nature of the pilot survey sample no formal analysis is done on the qualitative aspects of the responses, although this formal analysis will be done on the full survey to be done in the future. In addition, since this is a pilot survey detailed discussion on the implications of the results for the education program of actuaries in Australia will not be provided in this paper but left until the results of the full survey are available. As discussed in Section 3, the small nature of the pilot survey sample should induce caution in attempting to extrapolate the results of this survey to all actuarial graduates.

Finally, it should be noted that the survey responses did not indicate any significant issues with the structure of the survey. Answers to the open-ended questions indicate that respondents understood the questions being asked. The average time for completion of the survey was 19 minutes and the only comment from respondents on improvements in structure that was repeated more than once was a request to include a progress bar, which will be done in the full survey.

A. Part I

Figures 1 and 2 present the results of questions focused on Part I of the actuarial education program, the Core Technical (CT) subjects. The vast majority of respondents to the survey completed Part I via an exemption arrangement through an Australian university. The titles of the CT subjects are as follows; further information including syllabi can be found at the UK Institute and Faculty of Actuaries website:

CT1 – Financial Mathematics
CT2 – Finance and Financial Reporting
CT3 – Probability and Mathematical Statistics
CT4 – Models
CT5 – Contingencies
CT6 – Statistical Methods
CT7 – Business Economics
CT8 – Financial Economics

Figure 1 presents thoughts of the respondents on the usefulness of each of the CT subjects.

Interestingly, Figure 1 shows no clear differences in the relative usefulness of the CT subjects to respondents. This reflects the variety of fields the respondents work in and the breadth of the CT syllabi, with the usefulness of the CT subjects unsurprisingly tending to be correlated with employment area.
The only CT subject where more than 50% of the respondents answered 'A little bit useful' or 'Not useful at all' was CT5 – Contingencies, although of course those working in life insurance and superannuation regarded this as being 'Very useful'. Those working in non-traditional roles were more likely to find CT2 – Finance and Financial Reporting and CT7 – Business Economics useful.

Figure 2 presents overall thoughts of respondents on the strengths and weaknesses of Part I of the actuarial education program. In this and future figures the following list identifies the categories within which these strengths and weaknesses are measured:

- SC – Syllabus Content (i.e. the relevance of the actual content taught)
- SS – Syllabus Structure (i.e. the order in which the subjects and topics were taken)
- TQ – Teaching Quality
- TM - Teaching Materials (notes, online learning systems, etc.)
- A – Assessment
- F – Feedback on progress (from teachers and/or peers)

Figure 2 also shows no clear results, with all categories receiving at least 50% of responses at a 'Partial Strength' or higher level. According to these results, Part I appears to be broadly strong across all spectra. A number of open-ended comments referred to the wide variety in the quality of teaching. Other comments referred to self and peer development as being key aspects of Part I, for example:

It's really not about what we learned. It's about: (1) what you learn about yourself when faced with a tough academic challenge, and (2) the quality of the other students you interact with and learn from. Doing well in Actuarial instils a sense of self-confidence which cannot be matched by a finance or accounting degree.

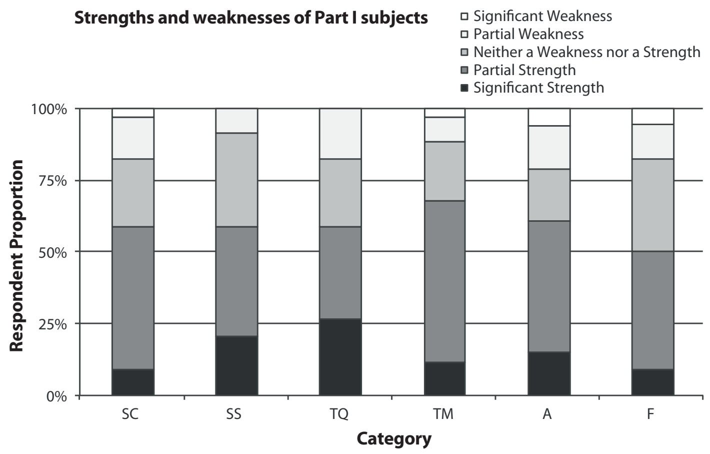

Figure 2: Rate the following as strengths and weaknesses of Part I in preparing you for your career.

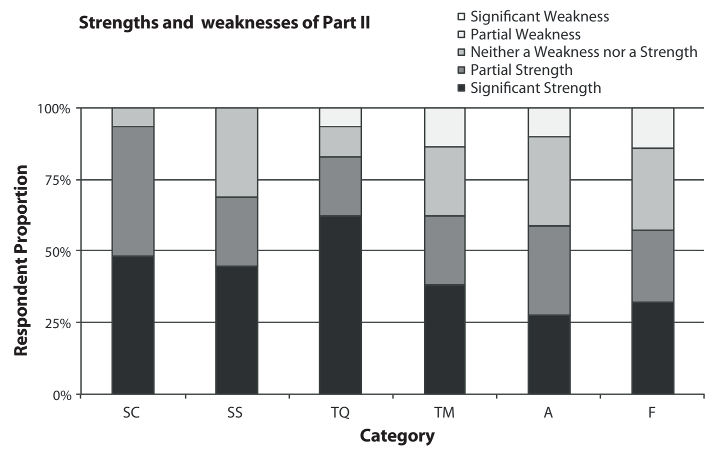

Figure 3: Rate the following as strengths and weaknesses of Part II in preparing you for your career.

B. Part II

Figure 3 presents the results of a question focused on Part II of the actuarial education program.

Figure 3 shows that Part II is considered to be stronger than Part I in preparing respondents for their career, with a higher proportion of respondents marking all categories except Teaching Material and Assessment at a 'Partial Strength' or higher level for Part II compared with Part I. This result is even clearer when it is noted that respondents who did not complete this question (most likely because they had not yet attempted Part II) tended to rate Part I as stronger on average than other respondents. The few open-ended comments tended to focus on the strength of the applied nature of Part II.

C. Part III

Figure 4 presents the results of a question focused on Part III of the actuarial education program.

Figure 4 shows a severe downturn in the quality of the education respondents perceive they have received in Part III compared to Part II. Respondents are generally happy with the content delivered in Part III, with over 60% rating Syllabus Content at a 'Partial Strength' or higher level. However, the lack of organised teaching and feedback is considered to be a weakness, with over 60% of each of Teaching Quality, Teaching Materials and Feedback receiving at least a 'Partial Weakness' or lower rating. Assessment was also considered to be on balance a weakness. This severity of this downturn is a potential concern to the profession, and it will be of great interest to see if it is replicated in the full survey.

Open-ended comments reflected these results:

There was no teaching or feedback mechanism.

The part III process seems to make a very limited attempt at actual education.

However, one respondent made specific mention of the benefits of the Australian-based Course 7A (Enterprise Risk Management) workshop in *understanding how ERM is applied in Australia. It also allowed for interactions that otherwise would not have been possible.* This respondent also stated that (in relation to the component of the course based on the UK ST9 course) *the materials were not structured very well - with many references to text books and areas of study overlapped.*

Whilst it is difficult to make significant inference from such a limited sample size, it appears that many respondents believe that Part III is more a hoop to be jumped through than a genuine attempt at education.

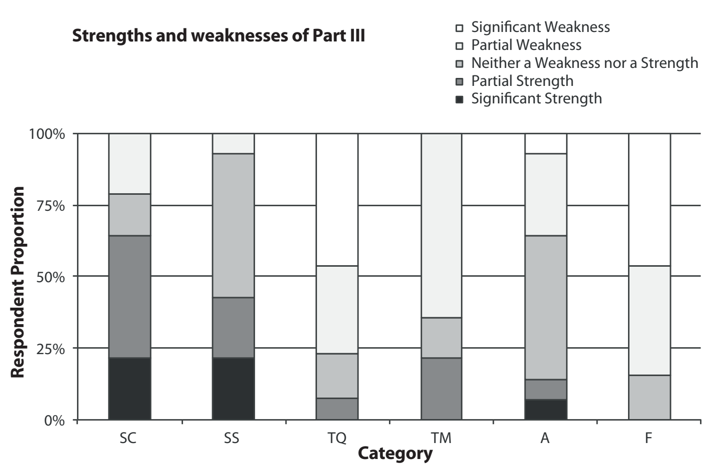

Figure 4: Rate the following as strengths and weaknesses of Part III in preparing you for your career.

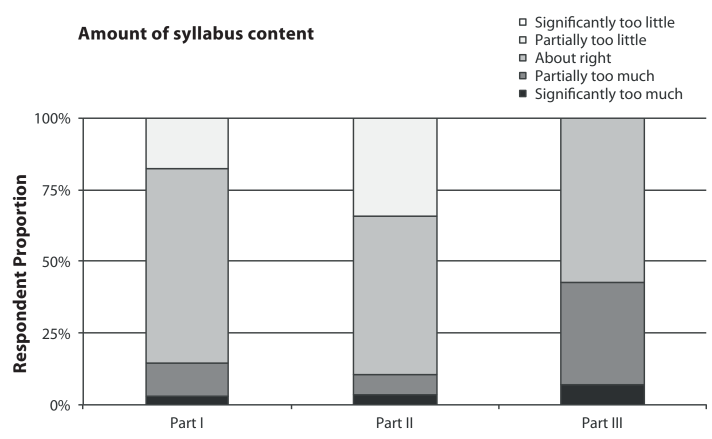

Figure 5: Comment on the amount of syllabus content taught in Parts I, II and III.

D. Other aspects of Parts I, II and III

Anecdotal evidence suggests concern amongst academics at the sheer amount of syllabus content required to be taught in Part I and Part II. Figure 5 presents the results of a question on respondents' perspectives on the amount of syllabus content across the three parts.

Figure 5 reveals that respondents did not appear to experience the syllabus overload that academics perceive in Part I and II. In fact, for all three parts more than 50% of respondents answered that the amount of syllabus content was 'About right', although respondents were more likely to perceive an overburdening amount of syllabus content in Part III than the other parts. Perhaps the anecdotal evidence described above is more a function of academics wishing to have space to teach additional material not in the syllabi than any significant overburdening of content in current courses taught by universities. It should also be noted that the nature of the selection of the pilot survey respondents is likely to have biased the results of this question towards a neutral response to this question for Parts I and II.

Whilst responses indicated the amount of Part I content was not overburdening, a number of the open-ended responses indicated concern at the focus of the content, in particular the lack of application of the content, for example:

Although the skills will be useful, the ability to apply them is more important.

I wish there were more focus towards practical applications rather than theory. For example, rather than spending too much time on Black Scholes model, it would be great to learn to price an option using a

Table 1: Where should the following non-technical capabilities be developed in actuarial education?

	CW	CO	L	PS	P	E	SB	T
Part I module	11%	7%	7%	21%	4%	4%	4%	11%
Part I integrated	46%	54%	11%	46%	25%	25%	18%	54%
Part II module	7%	4%	4%	7%	7%	7%	4%	4%
Part II integrated	57%	54%	36%	54%	50%	57%	43%	61%
Part III module	0%	7%	7%	11%	0%	0%	14%	4%
Part III integrated	29%	32%	29%	57%	36%	39%	46%	43%
Institute CPD	25%	25%	39%	14%	32%	32%	32%	36%

Note – in this question respondents were able to make multiple selections in each column.

standard software like excel or R that we would meet in real life (rather than an hypothetical example).

Other open-ended questions asked respondents to comment on the appropriateness of the relationship between Parts I, II and III. Many responses commented on the distinct differences between Part I and II, with some being positive and some being negative:

Positive

Part II is more practical compared to Part I. Part II is really useful to bridge the gap between tertiary education and the real commercial world.

Part II to me felt like a "How do you apply what you learnt in Part I?" for most parts of it.

I found the transition a welcome change from brute force mathematics, and Part II added a commercial flavour that definitely helped me transition into the workforce. Especially concepts like the actuarial control cycle, which can be applied in any discipline.

Negative

There was little relationship in that Part 1 was extremely quantitative while Part 2 was very much applied. Both elements are necessary although a clearer bridge linking them would have been useful.

There is a disjoint in the content taught in part I and part II. They're very different in nature, part I is quantitative whereas part II is very qualitative in nature. I don't think it is a weakness/strength, but it is interesting to see some students struggle with the former and do well in the latter and vice versa. Perhaps a more balanced approach will allow students with strengths in different areas to perform better on average in both part I and part II rather than to excel in only part I or part II.

E. Non-technical capabilities

We move now to the results of questions focused on non-technical capabilities, their importance to the careers of the respondents and the effectiveness and ideal location of their education, which are presented in Figures 6 and 7 and Table 1. The non-technical capabilities were drawn from the Actuaries Institute Capabilities Framework. Throughout the figures and tables that follow the list below identifies the capability on the axes/column titles:

CW – Communication (written) to a variety of stakeholders
CO – Communication (oral) to a variety of stakeholders
L – Leadership of people to a common goal
PS – Problem solving to provide practical solutions
P – Professionalism
E – Ethics
SB – Strategic foresight to anticipate business trends
T – Working in teams and managing relationships

The results from Figure 6 indicate that respondents attach a significant level of importance to the non-

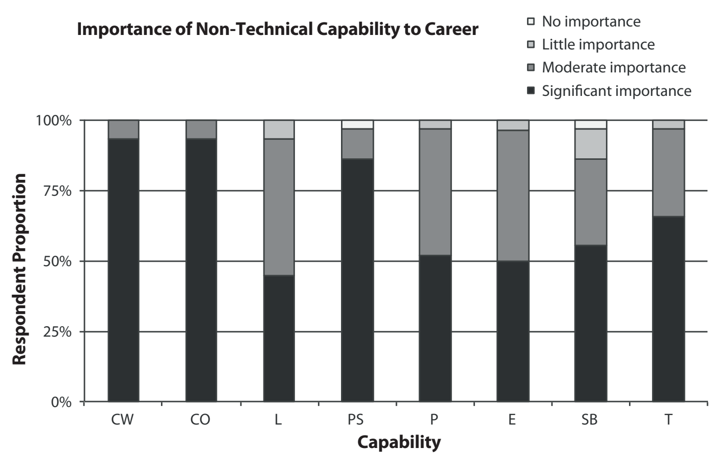

Figure 6: Rate the following non-technical capabilities in terms of their importance to your career.

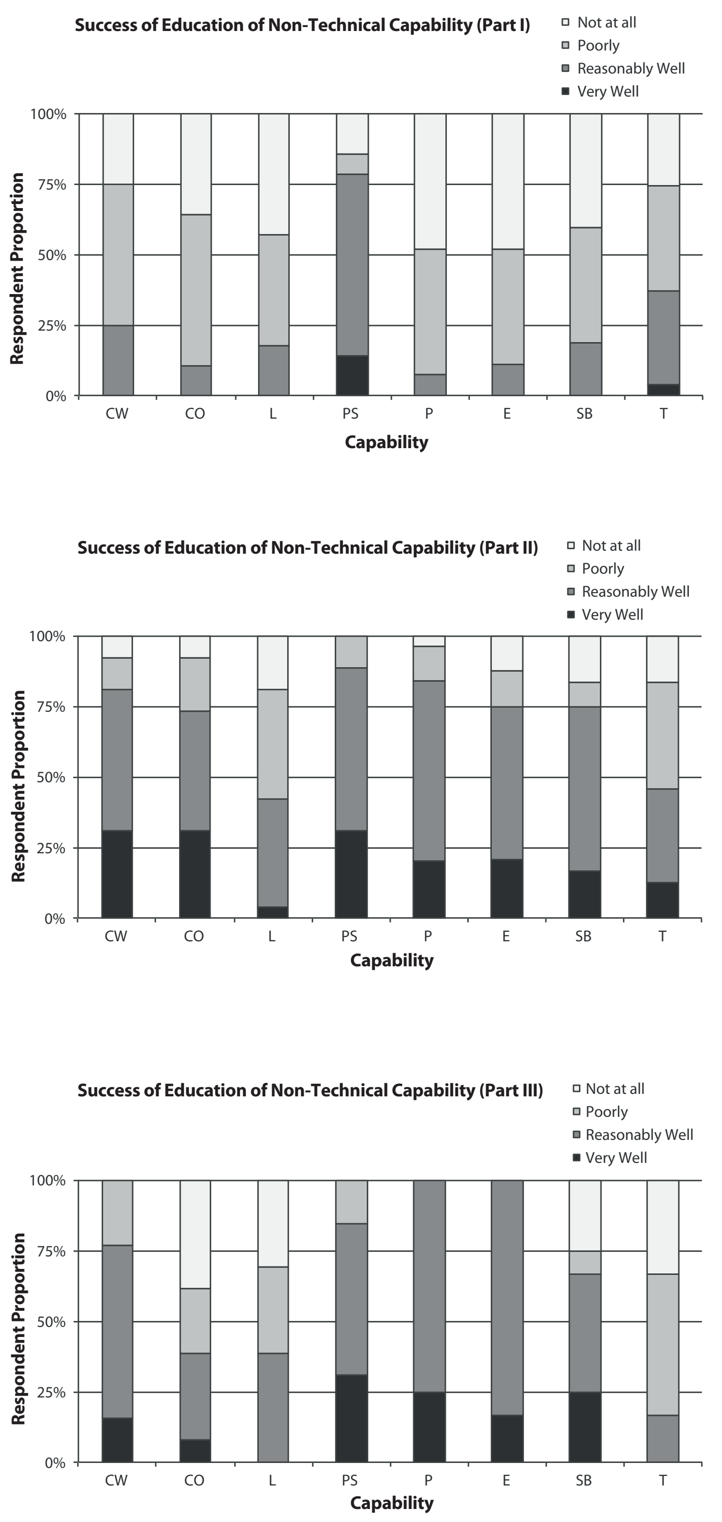

Figure 7 (I, II & III): How well were the following non-technical capabilities taught across your actuarial education?

technical capabilities described, with all categories, apart from Leadership, showing over 50% of responses having 'Significant importance', and Leadership being just under this figure. In all categories at least 85% of responses are at least of 'Moderate importance'. One respondent noted that *these skills are rated more highly than my technical skills at my work place.* Communication (both written and oral) appears to be valued the most by respondents, with over 90% of responses indicating 'Significant importance'. These results are consistent with the interview responses described in Butt et al. (2011).

Compared to the importance of non-technical capabilities described in Figure 6, Figure 7 paints a weaker picture of the quality of the education of non-technical skills across the actuarial program as a whole.

Part I is largely viewed as developing non-technical capabilities 'Poorly' or 'Not at all', with the exception of Problem Solving. It is interesting that even Teamwork is not considered to be developed particularly well within Part I, given that university business programs are widely known for using group work in teaching and assessment. Perhaps lecturers are reluctant to use group work in what they believe to be an overcrowded syllabus or alternatively they do not feel confident or equipped enough to use it. In any case these results are concerning given the substantial cohort of graduates of actuarial programs at Australian universities who will not go on to do Part II and III. Whilst these graduates may never become actuaries, their education in an undergraduate actuarial degree still has a significant impact on the reputation of the profession. Whether these results are due to poor program structure at universities or due to an overburdening of syllabus content enforced by the Institute is a debate outside the scope of the discussion of the pilot survey. Some argue (see Section 3) that university is not the appropriate location for the development of non-technical capabilities.

Part II fares much better in respect of non-technical capability development; it appears that respondents view the more practical nature of Part II as being more conducive to the development of non-technical capabilities. The majority of respondents still view Leadership and Teamwork as being developed 'Poorly' or 'Not at all' in Part II.

Responses for Part III are largely consistent with those of Part II, reflecting the relatively practical nature of Part III.

One potential reason for the inconsistency in responses between Figure 6 and 7 is that respondents view the acquisition of non-technical capabilities as something to be done informally over a period of work experience rather than or in addition to formal education. The results in Table 1 allow us to investigate where and how respondents believe that these non-technical capabilities should be developed. Where non-technical skills are to be developed in formal actuarial education, a clear preference is shown for integration into the curriculum rather than separate modules. Respondents are largely in favour of introducing Communication, Problem Solving and Teamwork at the Part I level, whilst Professionalism, Ethics and Strategic foresight are recommended to be introduced in later parts. The results for Leadership are not clearly in favour of any formal development, although a substantial number of respondents (39%) recommended development through Institute CPD.

Respondents were not given an option of ticking a box selecting none of the education options for any of the categories (which will be remedied in the full survey), although this can be inferred for each column by looking at the total number of people who selected at least one option in the entire question and subtracting the number of people who selected at least one option any given column. In all columns this turned out to be a very small number, with Professionalism at 11% showing the highest proportion of respondents not selecting a formal education option. This indicates that respondents are very supportive of non-technical capabilities being educated in formal programs rather than simply 'on the job'.

That said, respondents also noted the significance of on the job and other experience to the development of non-technical capabilities, with around 50% of respondents noting a variety of tasks that have contributed to their individual development.

F. Final comments

Respondents were asked a final open-ended question on their thoughts on the actuarial education program. A wide variety of responses was received here, with some respondents choosing to reiterate their previous answers and others choosing to make comments somewhat unrelated to the previous questions. Some made very positive comments about their education experience, for example:

My actuarial education prepared me quite well as I was able to think and logic things out better than some of my colleagues. The confidence developed also enabled me to push through on tough projects which colleagues with the same amount of work experience tended to be more hesitant in.

Others commented on the challenges facing the profession and potential responses through education programs, for example:

Actuarial education needs to first identify the unique differentiators that it has over other professions and then focus on developing these skills. Currently the education system is too focused on technical methods at the expense of practical commercial and personal skills development. As it was when I took the exams … it was training us to compete with computers on making impossible calculations rather than applying sound judgement and analytics to solve real world problems.

5. CONCLUSION

As noted earlier, there is a tension between what employers and graduates feel should be included in their education, particularly in relation to the generic skills required in the workplace, and academics who are responsible for delivering the education, who feel there is not enough solid theory in how these generic skills can be developed to warrant their inclusion in a university course. Perhaps also both parties take the view that helping people develop these skills can be very time consuming, so both are anxious to avoid that cost.

The graduates surveyed in this pilot study would like to see a greater component of generic skill development in the actuarial education process. In particular, communication and teamwork skills are considered important, and the pilot group of graduates feel these are best developed early in the education program. These results are similar to those found in surveys of accounting students.

The current syllabus for actuarial courses is considered by most academics that have taught these courses to be quite full already, and the graduates surveyed thought broadly that the technical content of the Part I courses was satisfactory in terms of its applicability to their employment, so the inclusion of a serious attempt to develop the generic skills would appear difficult without extending the actuarial program. While the pilot survey did discuss where respondents thought the generic skills should be developed, it did not discuss whether the respondents would trade off some technical skills in their Part I courses in order to have generic skills development included, and this is a question that could be included in future surveys. We would argue that a well-rounded education can prepare students to be lifelong learners in studying and applying technical skills that are not taught in Part I or other formal actuarial education.

While there were some comments on differing teaching standards in the Part I and II courses, generally the standard was considered significantly better than that in the Part III courses, which indicates that a change in the teaching and learning process for the Part III courses is needed, although the content was considered satisfactory and applicable to graduates work requirements. The survey does also indicate that there is a potential disconnect between the Part I and Part II courses, and perhaps some of the Part II content should be introduced at the Part I level.

The results presented in this paper are based on a small and non-random sample and have been published so that members, and in particular those involved in education review, can observe some of the preliminary findings of this research. Although the survey respondents were known to the researchers, the sample was deliberately selected to be what we believe to be reasonably representative of accredited actuarial university graduates in recent years. We intend to develop this work by conducting a substantially larger survey of recent graduates from accredited actuarial university programs and also of actuarial graduates from outside of the university accreditation system. We hope to obtain further feedback on our work and to encourage members of other actuarial professional bodies to engage in similar survey programs to further inform the development of actuarial curricula and education programs.

Bibliography

Butt, A., Chu, B. & Shepherd, J. (2011). Actuarial Education: Theory into Practice. Presented to the 46th Actuarial Research Conference, University of Connecticut, United States of America, August 11–13, 2011.

Chu, B., Evans, J. & Morgan, D. (2010). A Survey of Actuarial Graduates. *UNSW Australian School of Business Research Paper No. 2010ACTL17.*

De Lange, P., Jackling, B. & Gut, A.-M. (2006). Accounting graduates' perceptions of skills emphasis in undergraduate courses: an investigation from two Victorian universities. *Accounting & Finance*, 46: 365–386.

Howes, M. (2011). CEO's Column. *Actuary Australia*. March, page 30.

Ipsos Public Affairs (2010). 2009. Employer Study: Awareness and Perception of Actuaries Survey Findings. February 10, 2010.

Kavanagh, M.H. & Drennan, L. (2008). What skills and attributes does an accounting graduate need? Evidence from student perceptions and employer expectations. *Accounting & Finance*, 48: 279–300.

Keneley, M. & Jackling, B. (2011). The acquisition of generic skills of culturally-diverse student cohorts, *Accounting Education: An International Journal*, 20:6, 605-623.

Matthews, K.E. & Hodgson, Y. (2012). The Science Students Skills Inventory: Capturing graduate perceptions of their learning outcomes. *International Journal of Innovation in Science and Mathematics Education*, 20(1), 24–43, 2012.

Dr Adam Butt
adam.butt@anu.edu.au
Dr Adam Butt has been an academic at Australian National University since 2005. His research interests include individual choices relating to retirement and enhancement of actuarial education.

John Evans
john.evans@unsw.edu.au
John Evans is an Associate Professor of the Australian School of Business at the University of NSW. He is a member of the Council of the Institute of Actuaries of Australia, editor of the AJAP, Convenor of the Research Subcommittee of the Risk Management Practice Committee, and a Member of the Advisory Board of the Australian Centre for Commercial Mathematics.

Jim Farmer
jim.farmer@mq.edu.au
Jim Farmer has been an academic at Macquarie University since 1990. His research interests include life insurance discontinuance rates and the use of technology in teaching and learning.

Dr David Pitt
david.pitt@mq.edu.au
Dr David Pitt is an associate professor in actuarial studies. His main research interests are income protection insurance, statistical model selection with actuarial applications and actuarial education.

APPENDIX

The following pie charts present further information on the demographics of the 34 pilot survey respondents.

Year of last actuarial examination

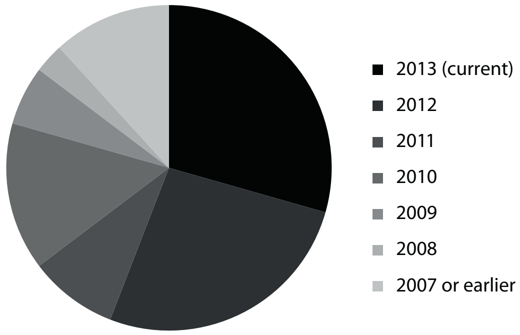

Area of employment

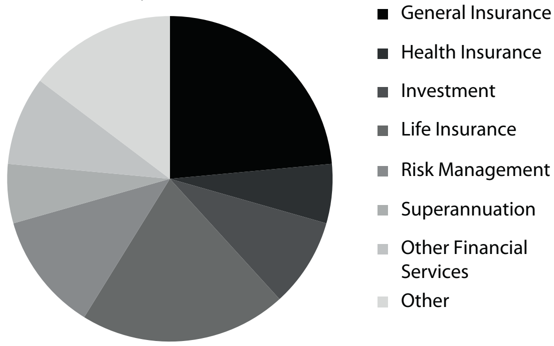

Location of employment

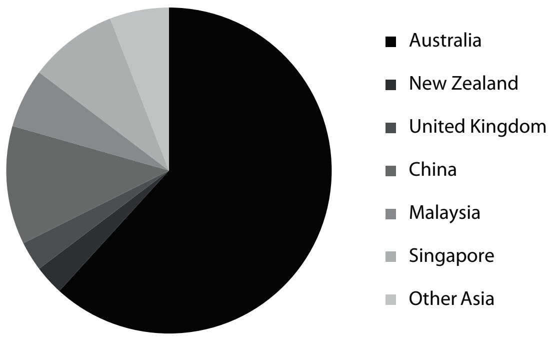

Years of full-time work experience

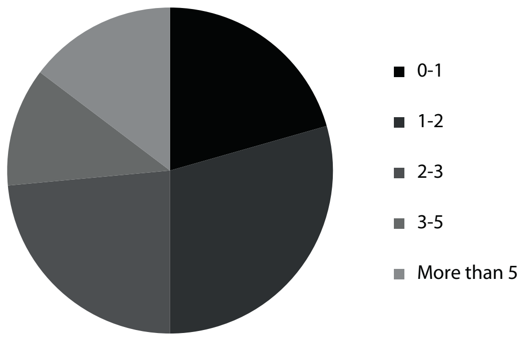

Age

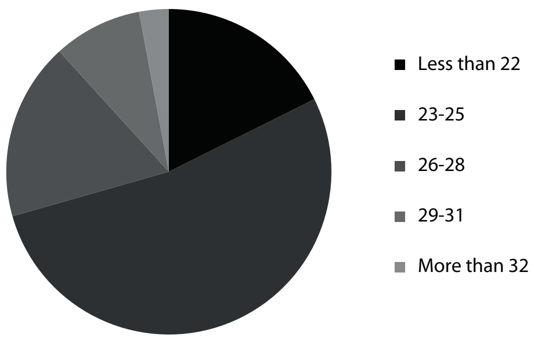

A survey of human capital risk management in Australian insurers

JOHN EVANS AND DR CAROL ROYAL

John Evans

Dr Carol Royal

Human capital risk permeates all other risks in a financial institution as was evidenced by the global financial crisis and subsequent events, but anecdotal evidence suggests this overriding risk was not well understood.

The authors applied for and were successful in obtaining a grant from the Actuaries Institute to survey Australian insurers as to how they identified and managed human capital risk. Discussions were held with 8 major insurers, a representative of the Actuaries Institute Risk Management Practice Committee and a representative of the prudential regulator, APRA. The general types of questions asked are set out in Appendix A.

The project identified that, while some aspects of human capital risk such as human resource management and occupational health and safety compliance, were well understood, the most important result of this survey was the identification that the linkages between human capital risks and other institutional risks was not well understood, which is why events like the LIBOR and fixed interest scandal in the UK could easily recur with significant reputational damage for the institution involved.

HUMAN CAPITAL RISK IN AN ERM PROCESS

Traditionally, risk management has evolved around risk identification, quantification, management, monitoring and review. This process has then lead to a situation where risk managers tend to identify and quantify risks within a static framework , but this process fails to be able to readily take into account the dynamism of risks, most of which it has been argued can exist simultaneously in the 'known', 'unknown' and 'unknowable' states .[1] The process of considering the statistical attributes of a risk simply captures the current weighting to the various states and is an inherently unstable process. A more stable process is to identify the drivers of risk in an institution, and then develop an appreciation of how these drivers influence the various risks that may then arise. By concentrating on the

1 Forthcoming article 'Dynamisms of risks and its risk management implications' (ed). *Risk Management Today*, Corrigan J., Evans J., and Ganegoda A.

KEYWORDS
human capital, enterprise risk management

AJAP 2014; 1: 77-80

drivers, and how they are changing will lead to a better understanding of the likely risk events that could occur and their importance and interrelationships. To state the obvious, it is future risk events that concern an institution, not the historical ones, and this is well illustrated in the paper by Josh Corrigan and Neil Allan,[2] where they looked at the drivers for major operational risk events and found a group of common drivers, which if evident in an institution would likely be indicative of the occurrence of a major operational risk event.

Figure 1 illustrates our view of the relationship between the major drivers of a risk event in an insurance institution, and the risks.

In Figure 1, 'impact' relates to the effect on the balance sheet of the institution, and the risk 'level' refers to the relationship between the risk events.

By considering the drivers of risk events, it becomes clear that a significant number are related to human capital risk, which for our purpose is defined as dysfunctional people management systems associated with recruiting, training and development, performance management, remuneration adherence to ethical standards and succession planning. Human capital risk impacts, but is not necessarily the sole driver of all, of the risks identified as 'Level 2' risks for an insurer and is therefore considered a major risk for an institution requiring significant attention in the risk management process.

2 'Emerging Risk Assessment-Latest Innovation and Practice', Allan N & Corrigan J, presented to Actuaries Institute Summit, May 2013

Figure 1: the relationship between major drivers of a risk event in an insurance institution and risks.

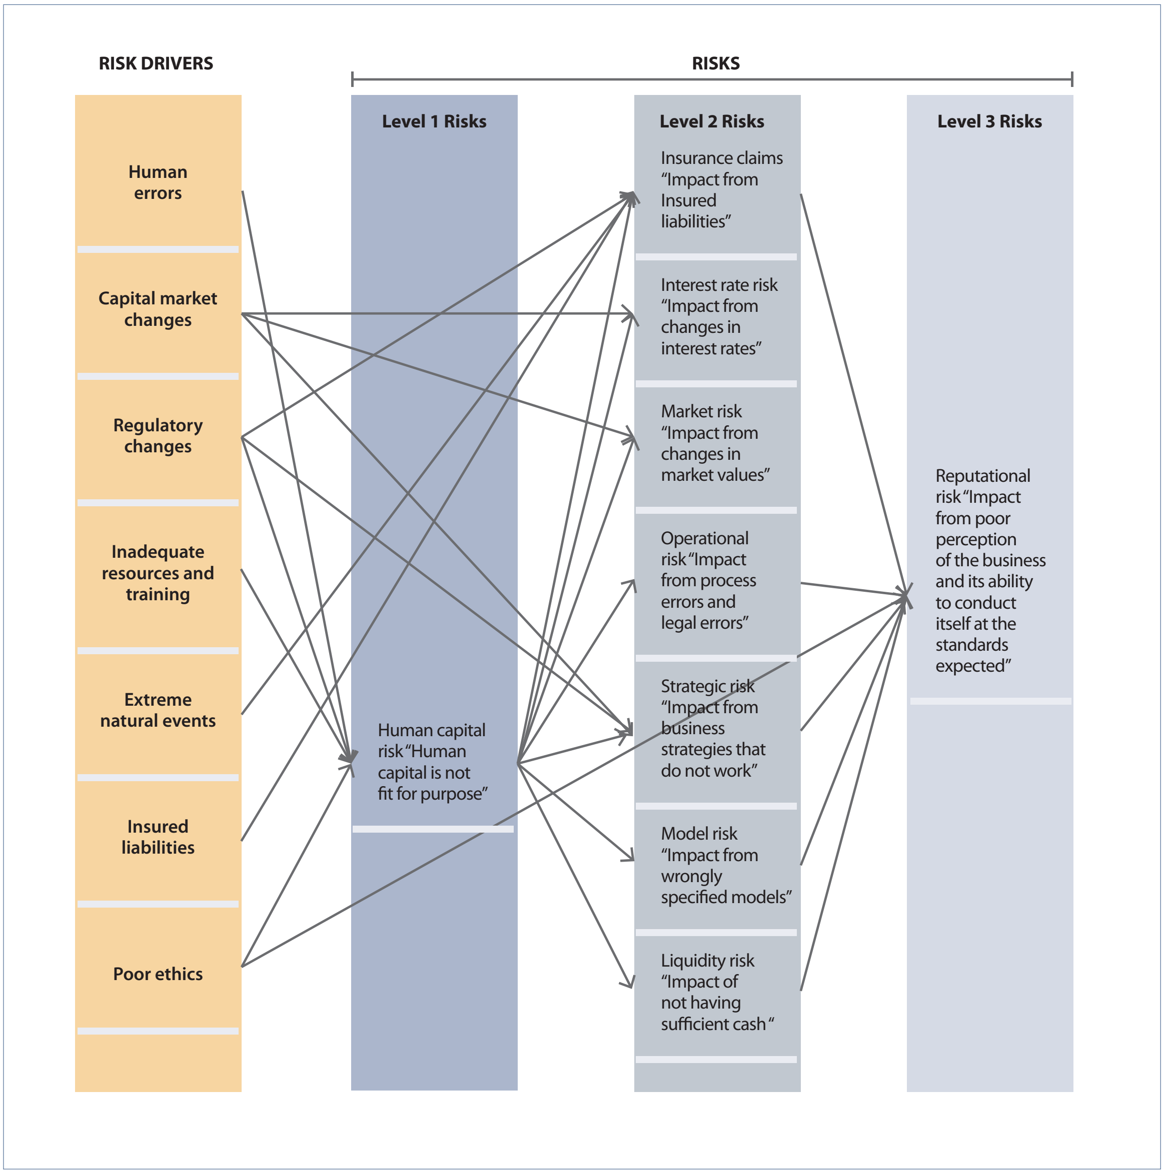

SURVEY RESULTS

Linkage of human capital risk and the ERM process

The survey highlighted that in general there was no specific knowledge, expertise or understanding of the role of human capital in the ERM process. Human capital risk was more often than not confused with human resource activities such as recruitment, performance management, remuneration management and some understanding of culture issues. Participants adopted an internal company-focused lens and therefore indicated that all these human resource activities appeared to operate within insurance organisations from an organisational functioning perspective and in an internal world. There was no evidence of the link between human resources activities and the stated enterprise risk management framework.

Institutional knowledge base

When asked about the knowledge base/qualifications and skill base of the Board, CEO, CFO, compliance officers and any other staff whose roles are focused on compliance and enterprise risk management in their organisations, there was no evidence of any of the organisations and their board members having the appropriate skill sets in human capital risk. Nor was there any indication of human capital risk knowledge skills or expertise at the more internal senior levels below the boards.

Training and development

The participants indicated that training and development in human capital risk management was patchy. Training and development mainly concerned ongoing development with regards to the traditional organisational HR activities pertaining to leadership, communication, transparency and governance/ethics which took for the most part the form of workshops for senior staff, being organised more around compliance than risk management. Some of the themes, such as risk culture, and ethics/whistle blowing, were captured in training sessions for all staff as well. Human capital risk was seen by the organisations to be more generally related to behaviour and its link to compliance within their organisations.

Some insurance organisations use culture and staff engagement surveys to ascertain the level of risk management culture and appetite for risk. These human resource surveys are used for motivating and retaining staff for the most part and for providing a platform by which staff can raise issues that are concerning them. Some of these human resource surveys used by the insurance companies cover a couple of questions on risk management. Others use focus groups to understand staff concerns more thoroughly and try to glean whether there is anything troubling in those surveys that relates to the behavioural aspects or culture of their organisations that could lead to a risk situation. However, most organisations said that they accepted that staff are internally focused about their thinking and knowledge of risk-including human capital risk.

Reporting with respect to risk management, especially human capital.

Most of the participants indicated that the reports produced by their organisation for APRA did not encompass human capital data. However, some participants noted that they held discussions with APRA on the risk management framework, prudential and remuneration reviews, which included human capital risk aspects. Most companies produced internal reports for audit and governance committees with respect to risks, sustainability and short-term incentive targets. Some of these key internal reports were reviewed and signed by the boards, and most companies declared that they undertook assessment of risk frameworks before their board signed the declarations and annual reports given that APRA held the board responsible with regard to risk and mitigation. However, it was noted that most reporting was driven by regulatory requirements. The human capital risk focus within reports produced for APRA was therefore limited to specific issues such as 'fit and proper' test on key staff or availability of adequate human resources for the proper functioning of the organisation.

The potential for inconsistency between the rhetoric and reality

Most participants were quick to agree that inconsistency existed across the industry between the rhetoric espoused at the top levels of the organisation as opposed to the reality of what is practised within organisations with regard to human capital risk. Although the majority of the participants thought that organisations made genuine efforts and took action to reduce risk, most pointed to potential misinterpretation of strategies affecting company culture as the source of such disconnect.

Appreciation of the value of human capital risk management

Some saw successful human capital risk management as having a commercially competitive advantage as well and were keen to develop their organisation's engagement with the issue, but the majority were confused as to how to define and manage human capital risk.

CONCLUSION

To a large extent human capital risk is being managed as a risk by itself and not as a risk that drives other risks. The survey showed that there was confusion by the participants on how to identify, manage and integrate human capital risk management into the ERM process. This appeared to stem from a limited understanding of general management systems thinking which looks at systems, be it risk or other systems, in terms of inputs, processes and outputs.

The survey has shown that without a change in the approach to managing risks in general, ie considering drivers of risks as the starting point of risk analysis, the industry remains exposed to unexpected detrimental events, particularly related to failed human capital systems.

APPENDIX A

The survey was carried out by discussion with the interviewees, and whilst the discussions then varied, the following were generally discussed:

1. How do you achieve and monitor whether a risk management framework is being achieved in your institution, and where change is required in the risk management framework, how do you monitor the change process?
2. What is the knowledge base of the Board, CEO, CFO, compliance officer/s and those whose roles are focused on compliance and enterprise risk management in your organisation?
3. What is the system of recruitment for those staff where risk management is key to the role?
4. What are the key performance management systems and KPIs (key performance indicators) expected of staff where risk management is key to their role?
5. What kinds of reports does your organisation produce with respect to risk management, especially human capital?
6. Does your staff have adequate knowledge of the human capital systems such as risk management culture, staff engagement levels, change management, that are necessary and need to be evident to support an overall enterprise risk management framework in your organisation?
7. What elements of human capital risk are managed within your organisation?
8. From a human capital perspective, is there the potential for inconsistency between the rhetoric and the reality provided by financial institutions in the sector?
9. Consider recent financial services scandals, such as LIBOR and Barclays Bank. What, in your opinion, are the key causes of such situations?

John Evans
john.evans@unsw.edu.au
Associate Professor John Evans, of the Australian School of Business at the University of NSW, was the initial Head of the School of Risk and Actuarial Studies from 2009 to 2012. He is a member of the Council of the Institute of Actuaries of Australia, editor of the Australian Journal of Actuarial Practice, Convenor of the Research Subcommittee of the Risk Management Practice Committee, and a Member of the Advisory Board of the Australian Centre for Commercial Mathematics. He chairs risk and compliance committees for several institutions. Previously, amongst other roles, he was a Guardian of the New Zealand Superannuation Fund. He was also the inaugural Academic Director of the Master of Finance Program for the School of Banking and Finance at UNSW.

Dr Carol Royal
c.royal@unsw.edu.au
Dr Carol Royal is Director of the Masters of Technology Management at the Australian School of Business at the University of New South Wales and Honorary Visiting Fellow Cass Business School City University London. Her research evaluates publicly listed companies through human capital analysis and its links to corporate performance, corporate sustainability and change management systems, using qualitative research methodology. She provides consultancy services to enhance the quality of intangibles analysis in the equity market, having previously worked in industry as a senior human resource manager.

The fallacy of using superannuation taxation expenditures

DR DAVID KNOX

Dr David Knox

BACKGROUND

The taxation of superannuation has been subject to much speculation and commentary in recent times. This discussion is prompted every year when the Commonwealth Treasury produces the Tax Expenditures Statement, which provides details of more than 350 tax expenditures. The latest Statement, published in January 2013, showed the largest projected tax expenditures for 2012–13 are for superannuation and owner-occupied housing, with the superannuation tax expenditure shown as $31.9 billion with an expected growth to $44.8 billion in 2015–16.

Naturally such significant numbers attract considerable attention from interested parties who may be aiming to increase government revenue or improve equity within the community.

This paper will show that the Treasury approach, which ignores Government savings from reduced Age Pension costs, is fundamentally flawed and, if used to develop long-term retirement income policy, is likely to lead to sub-optimal outcomes for individuals, households and the government.

A WORKED EXAMPLE

Initially let us consider Treasury's approach to the calculation of the tax expenditures for superannuation. There are several items that are included in the total superannuation tax expenditure, but two of them make up 95% of the total expenditure:

1. the concessional taxation of employer contributions
2. the concessional taxation of superannuation fund earnings.

For most Australian workers their employer superannuation contributions (including the compulsory 9.25% Superannuation Guarantee) are taxed at 15% and not at the individual's marginal tax rate. Similarly, the investment income earned by their superannuation fund is taxed at 15%, which is often reduced to a net effect of about 8% after allowing for dividend imputation and capital gains tax credits.

KEYWORDS
superannuation, taxation, age pension

AJAP 2014; 1: 81-84

Figure 1 shows the annual value of these two taxation concessions (as calculated by Treasury) for an average income earner[1] over an assumed working career of 40 years. In effect, Treasury uses the difference between the individual's marginal tax rate (i.e. 34% for average income earners, including the Medicare levy) with the 15% superannuation tax rate.

The concession in respect of the contributions gradually increases as the average wage is assumed to increase by 4% p.a. However, the concession in respect of the investment earnings rises more steeply as the member's balance in the superannuation fund increases with both contributions and investment earnings. This means that, in the latter years, the concession in respect of the investment earnings becomes more valuable. This is consistent with the Treasury's numbers, where the investment earnings concession is now 30% higher than that for contributions.

However, concentration on the working years is only part of the story for an individual's retirement arrangements. The provision of retirement income is a journey through both the accumulation years and the payout phase; both must be considered to assess the full picture.

Figure 2 shows the savings to government of the future age pension payments in respect of an average income earner who works for 40 years and then has a retirement period of 20 years. It is clear that the accumulated superannuation benefit reduces future age pension payments at a significant saving to government. These savings are totally ignored in the Treasury figures. Even the savings shown in the graph are not as

1 It is assumed that this individual works full time for 40 years and earns the average wage (currently about $72,500) each year. The assumptions used in the modelling are outlined in the Appendix.

Figure 1: The calculated costs for an average income earner.

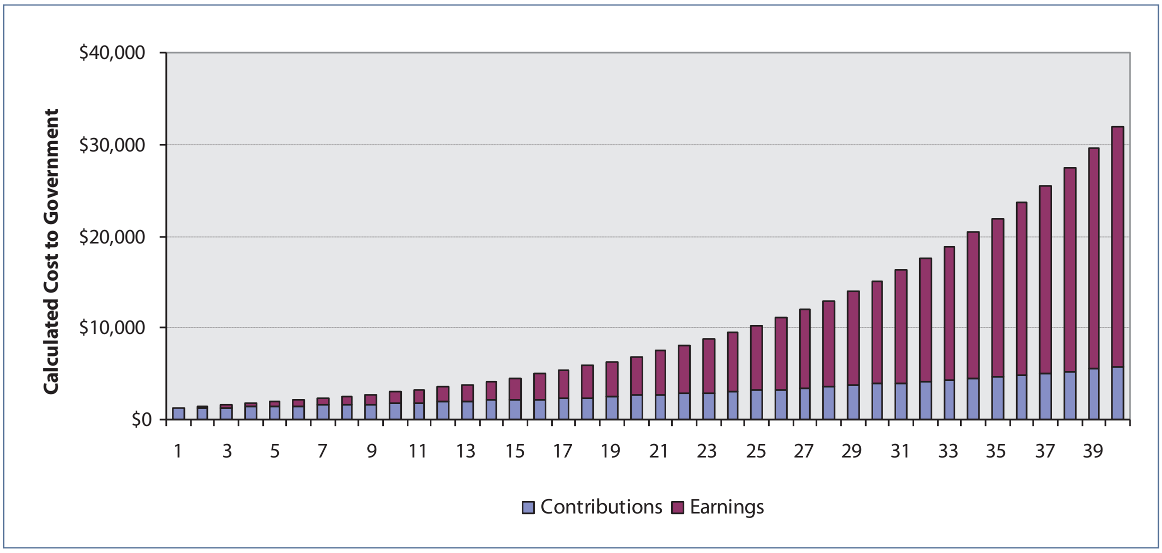

Figure 2: The calculated costs and savings for an average income earner.

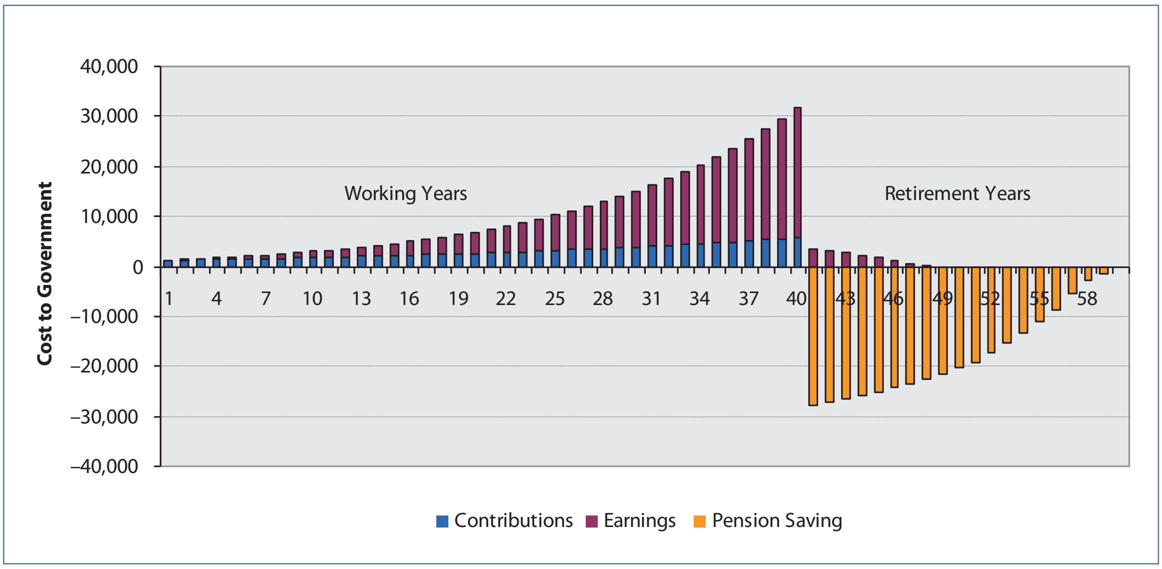

high as they could be, as it is assumed that 15% of the superannuation benefit will be used immediately at retirement.

It is also noted that these savings reduce each year as the retiree is assumed to gradually draw on their account based pension for income in line with the government rules. This is consistent with experience as many retirees commence with a part Age Pension but receive a full age pension in their later years.

There are also further concessions in respect of investment earnings in the retirement years. However, these are much lower in retirement as the individual's income and marginal tax rate have reduced.

The above example considered an individual earning the average wage throughout their working career. Table 1 considers the present value of the taxation concessions and Age Pension savings at different lifetime income levels across their lifespan.

Table 1 highlights the following findings:

- the value of the taxation concessions rises with income, which is to be expected as both the amount of contributions and marginal tax rates rise with income
- the Age Pension savings also increase with income as the superannuation balances are higher thereby leading to greater Age Pension savings in the future
- the net cost to government, after allowing for the Age Pension saving, reduces as a percentage of the calculated concessions, as income rises.

It should also be noted that the concession on investment earnings, as calculated by Treasury and shown in the above table, is overstated for most individuals. The reason is that most individual taxpayers do not pay tax at their full marginal rate on their non-superannuation investment income. The reasons include the availability of imputation credits and capital gains tax concessions; the opportunity to invest through a lower income partner; geared investment opportunities (such as negatively geared property); and the common practice of investing in the tax-exempt family home.

The range of incomes shown in the table cover the vast majority of full time workers for the following reasons:

- A low income earner (at say 0.5 times the average wage) is likely to represent part time or casual workers. If these individuals continued to receive income at these levels throughout their career, they would receive some taxation concessions but would also receive the full Age Pension due to their small superannuation benefit.
- Whilst some individuals earn above twice the average wage, there are very few employees who have earnings above this level throughout their 40 year career. It is critical to assess superannuation throughout an individual's career and not on a year-by-year basis.

Table 1: The present value of the concessions and Age Pension saving at different income levels assuming a 9% SG contribution.

Income level (multiple of average wage)	0.75	1.0	1.5	2.0
Concession on contributions	35,762	47,683	88,464	117,952
Concession on investment earnings	67,903	93,347	195,091	288,032
Total concessions	**103,665**	**141,030**	**283,555**	**405,984**
Age Pension savings	–13,251	–54,383	–151,448	–223,073
Net cost	**90,414**	**88,647**	**132,917**	**162,911**
Net cost as a % of tax concessions	**87.2%**	**62.9%**	**46.9%**	**40.1%**

ADDITIONAL CONCERNS

There are several other reasons why tax concession calculations should be treated with caution when considering changes to superannuation policy, including the following:

- No allowance is made for behavioural change. It is inevitable that if the taxation concessions for superannuation are reduced, Australians will change their investment and look for other tax-effective investments. The status quo will not remain.
- There is an element of double counting. If future contributions were to be taxed more heavily there would be reduced investment income in the future, thereby reducing the value of that concession.
- The level of the concessions is directly influenced by the level of marginal tax rates. This means that the recent reform to increase the tax-free threshold actually increased the value of the superannuation concessions for middle income earners, with no change to superannuation.
- The Tax Expenditures Statement represents a one-year snapshot of a long-term investment. This approach may be reasonable for a tax expenditure item that is short term, such as the GST exemption on fresh food, but it is inappropriate for the development of a long-term sustainable policy as Australia prepares to cope with the pressures of an ageing population.

One final misconception needs to be highlighted. It is commonly assumed that should all the superannuation tax concessions be removed, then the budget revenue would increase by the level of these tax expenditures. This is not the case, as individual behaviour would change and the potential revenue gain would be much lower than the quoted value of the tax expenditure.

CONCLUSIONS

The taxation treatment of superannuation is an important long-term policy for both the Australian economy and for Australians as they prepare for a dignified retirement. However, superannuation is only one part of Australia's well-respected multi-pillar retirement income system. A more holistic approach, taking into account both superannuation and the Age Pension, is needed to shape the long-term policy.

The Henry Tax Review supported the concept that superannuation should continue to be taxed more favourably than other forms of saving. The Tax Expenditure Statement makes no mention of this policy, nor does it highlight that the concessional treatment of superannuation is less generous in Australia than the policies adopted in many other developed economies.[2]

Increasing the taxation of superannuation has two obvious effects – it reduces future superannuation benefits and thereby increases future Age Pension payments. Such an impact does not just affect low income earners. It will affect middle and higher income earners; many of whom will receive a part Age Pension in future years.

Whilst increasing the taxation of superannuation may be a quick fix to respond to an immediate revenue shortfall, a much better long-term solution would be to improve the integration between the superannuation system and the means-tested Age Pension. This could be achieved through a revision of the means testing arrangements or through the encouragement of income streams from superannuation, possibly through restricting larger lump sum payments. Such approaches would represent a long-term sustainable approach and are much better than further tinkering with the tax system, which will also reduce the community's confidence in Australia's retirement income system.

2 This result was highlighted in Mercer's publication entitled *Tax and Superannuation: Benchmarking Australia against the world's best retirement savings system,* which was released in February 2013.

APPENDIX

Modelling assumptions

The primary purpose of this research is to consider the level of government support for an individual's superannuation over their lifetime and the impact that this will have on their future age pension payments.

The underlying assumptions used in the calculations are described below.

Net investment earning rate (after fees and taxes)

- Accumulation period (pre-age 65) — 7% p.a.
- Post retirement period — 6.5% p.a.[3]

Discount rate — 4% p.a.

This rate was chosen as it reflects the expected growth of average wages over the longer term, representing a combination of inflation and productivity increases.

Age Pension level for a single individual at March 2013 — $20,667 p.a.

Income tax scales

- As applying for 2012–13 with the marginal tax thresholds indexed at 4% each year.

3 Although the account based pension pays no tax, a slightly lower rate of investment earnings has been assumed due to the higher level of conservatism adopted by many retirees.

Dr David Knox
david.knox@mercer.com
Dr David Knox is a Senior Partner at Mercer and author of the Melbourne Mercer Global Pension Index. He has been an independent Board member of APRA and President of the Institute.

Long-term personal injury insurance and the NDIS

JOHN WALSH AM

John Walsh AM

KEYWORDS

NDIS, National disability insurance scheme, disability reform, disability funding, social welfare reform, social welfare funding

AJAP 2014; 1: 85-86

The actuarial profession has made a significant contribution to the prudential governance of long-tailed statutory insurance classes for the past 30 years. Our work in general insurance business began in earnest during the 1970s and accelerated after the work of Sir Owen Woodhouse introduced a new way of thinking about compulsory workers compensation and personal injury insurance in general.

The Woodhouse principles focused very much on equity and durable outcomes for the injured person rather than common law retribution against negligent parties, usually in the form of lump sum monetary settlement. The Woodhouse principles are:

Community responsibility
Comprehensive entitlement
Complete rehabilitation
Real compensation
Administrative efficiency

This thinking transformed the nature of many of Australia's accident compensation schemes during the 1980s, leading to longer term entitlement to weekly benefits and treatment support, and forcing actuaries to take a much longer tailed modelling approach to statutory liabilities. The notion of 'the scheme actuary' emerged for workers compensation and some transport accident schemes.

As these schemes matured it became obvious that short-term claims management targeted at the bulk of injuries was often not appropriate for people with more complex injuries. In particular the most severe injuries triggered the idea of special management teams. The Victorian Transport Accident Commission established its major injury division, and the NZ ACC established a national serious injury service. More generally, Australia's insurance ministers called for a report into a national long term care scheme for 'catastrophic injury', and New South Wales continued this work in establishing its Lifetime Care and Support Scheme in 2006 – which covers predominantly people who sustain major spinal cord injury and brain injury in a motor vehicle accident, and need lifetime support for basic activities of daily living.

From this model it was not a huge leap to develop a similar approach to the lifetime support planning and service delivery for people with all types of disability – this thinking underpins the National Disability Insurance Scheme (NDIS).

From an actuarial point of view, the so-called 'insurance approach' comprises the following principles, all of which were new to traditional social welfare funding and service delivery:

Risk pooling: the Woodhouse notion of community responsibility, clearly asking the Australian population to contribute to the risk of acquiring a significant disability with lifetime support needs, and to support those who already have a disability

Actuarial cost estimation: the expected cost of providing this care and support is estimated from the ground up using a scientific actuarial approach of prevalence (claim frequency), and average support need (claim size) of the target population.

Reasonable and necessary entitlement: the 'entitlement' of an individual participant is determined based on their support needs, rather than on available funds in a rationed system.

Active outcome management: participants entering the NDIS are invited to think about their goals and aspirations within the limit of their reasonable and necessary resource allocation, with an expectation that active community participation will provide a positive cost benefit to the scheme.

Independent prudential governance: it was recommended by the Productivity Commission that the NDIS be operated by an independent statutory agency with an independent Board selected based on their experience in the areas of commerce, insurance and disability. While the NDIS is necessarily a pay-as-you-go scheme, the discounted net cost of future expenditure will be estimated by the scheme actuary and monitored by the Board.

Longitudinal real-time data and reporting by a scheme actuary: the NDIS will collect longitudinal information on participants regarding support needs, service utilisation and cost, and participant outcomes measured against their individual plans. The scheme actuary will translate these data into a quarterly monitoring report and annual financial condition report to the Board. The Board will be accountable to the relevant ministers and committees, and ultimately to the Australian public.

Effectively, it is anticipated that the NDIS will operate in a similar function to a statutory accident insurance scheme, but with notional balance sheet reporting and investments. Premium contributions from policyholders will be simulated by a 'group contribution' from Australian governments – in a manner not dissimilar to group life cover by an employer in a superannuation scheme. The challenge for the NDIS, its Board and its actuary will be to have the information, strength and influence to maintain the principles of the scheme in changing political and economic environments.

John Walsh AM FIAA
magoo.actuarial@gmail.com
John Walsh AM has recently retired after 20 years as a partner at PricewaterhouseCoopers. He has been adviser to many of Australia's accident compensation schemes and health and disability systems. During 2010–2011 he was Associate Commissioner of the Productivity Commission's inquiry into Disability Care and Support, which designed and recommended the National Disability Insurance Scheme.

John is a member of the Board of the National Disability Insurance Agency. However, the views in this note are totally John's, and not necessarily those of NDIA or any of its officers or other Board members.

Stress testing for insurers

LEN ELIKHIS

Len Elikhis

Prudential Standards LPS and GPS 110 (Capital Adequacy) have positioned the ICAAP at the centre of the insurer's capital management framework. As part of the ICAAP, insurers are required to undertake stress testing of their risk exposures to evaluate the impacts of exceptional but plausible events. Within an insurance context these techniques include sensitivity analysis, scenario analysis, and reverse stress testing. In this note I discuss how insurers can obtain value from the stress testing exercise.

Many capital models employ the correlation matrix to linearly aggregate risk exposures. While the approach simplifies calculations and makes the models tractable, it fails to address the non-linear relationship between risks and capital within the tail of the distribution. In other words, the impact of multiple risks occurring simultaneously is greater than the sum of the risks after allowing for diversification. Furthermore, risk interactions within the tail are non-constant and can be influenced by market forces (for example, anti-selection spirals and market panics). These complexities are beyond the scope of the standard capital model and stress testing provides a means by which most conceivable risk interactions can be directly specified and analysed.

Stress testing can also aid in the iterative process of articulating the Risk Appetite Statement and evaluating the appropriateness of the target surplus position. In his article titled 'Calculating non-linearity in an ICA: the killer scenario', Peter Kettlewell (2007) demonstrates how principal component analysis can be used to generate stress scenarios at different levels of confidence using an existing capital model and the risk correlation matrix. The technique can be used to mathematically produce scenarios that are within (or beyond) the Board's Risk Appetite and these scenarios can be incorporated into narratives that express the boundaries of the Risk Appetite Statement in less conceptual terms.

Stress testing can be extended beyond a mathematical calculation into a simulated risk management exercise that helps to validate aspects of the ICAAP and improve preparedness for adverse events. For example, detailed consideration of an adverse scenario may help to address the following questions:

- Are the capital triggers appropriate?

KEYWORDS

stress testing, ICAAP, capital, risk management

AJAP 2014; 1: 87-88

- Is there sufficient opportunity to restore the capital position before a breach of the PCR?
- Are assumed risk mitigants (reinsurance, hedging instruments, etc.) effective and available on economic terms?
- Are assumed actions to restore the capital base realistic under a stressed scenario?
- Can the capital position be estimated on a reasonably frequent basis?

The stress testing exercise can also be used to consider those scenarios which aren't adequately addressed by capital models (operational risks, such as fraud) or for which capital isn't an effective mitigant (legislative risks, such as legal constraints on choice of underwriting factors).

But can stress testing, poorly executed, lead us astray?

Following a review of the internal stress scenarios adopted by Australian ADIs, APRA observed that 'too often, the scenarios and the modelled results are not severe enough to result in losses. This disaster myopia needs to be overcome.' ('Stress-testing for authorised deposit-taking institutions', 2010). In other words, financial institutions need to seriously and dispassionately contemplate those scenarios that may undermine their capacity to remain solvent. If we step outside the financial context and look at the lessons from the recent Fukushima nuclear disaster we can gain a useful perspective into the nature of disaster myopia.

In March 2011, a magnitude 9.0 earthquake produced a tsunami that struck the Fukushima Daiichi nuclear power plant, leading to flooding which eventually led to a core meltdown at three reactors. Following the incident it emerged that the operator of the plant, TEPCO, had chosen to disregard an internal stress test that exposed the plant's vulnerability to tsunami waves exceeding 10 metres and recommended immediate works to improve protection from flooding seawater. The study noted that a magnitude 8.0 earthquake was strong enough to produce a 10-metre wave and cited seismic evidence from as recently as 1896 and 1933 in the geographic area. Unfortunately, TEPCO's management disputed the modelling and dismissed as remote the possibility of a 10-metre wave, even under a magnitude 9.0 earthquake scenario.

So what might we do to try to reduce our own risk myopia?

Firstly, we need to acknowledge that we are constrained by a paucity of stable data. Unlike seismic activity, which develops on a geological scale, the behaviour of financial markets and insurance risks evolve far more rapidly. The impacts of financial innovation, globalisation, societal change and medical developments are difficult to predict and can introduce material discontinuity between past and future trends.

We should also remain sceptical of the predictive power of capital models. Unlike seismologists who deal with a single causal relationship – that is, between an earthquake and a tsunami – financial models need to contemplate dozens of uncertain and dynamic interactions, including the risk of systemic failure and contagion. In some cases, the impact of an extreme event may result in secondary effects. Some of these effects may directly follow from the scenario, while others may arise from systemic impacts. Insurers should aim to consider the totality of effects in order to arrive at a realistic estimate of the financial state after an event. Examples of secondary effects include:

- limited or no access to additional reinsurance
- partial reinsurance recoveries
- impact of policyholder behaviour
- impact of rating down-grades
- impact on capital costs
- competitor behaviour
- adverse feedback loops.

Consequently, whilst stress testing cannot precisely capture the nuances of real events, it offers significant value in terms of modelling tail risk, communicating and refining the Risk Appetite Statement, and improving general preparedness for adverse events.

Bibliography

APRA (2010). Stress-testing for authorised deposit-taking institutions. *Insight*, p. 2–12 .

Kettlewell, P. (2007, November). Calculating non-linearity in an ICA: the killer scenario. *Insurance and Financial Services* Review, p. 15–18.

Len Elikhis FIAA
len.elikhis@nab.com.au
Len Elikhis is an actuary working within MLC's Market Risk Products team. Prior to joining MLC, he was with Deloitte Actuaries and Consultants, where he assisted a number of life companies to implement LAGIC.