

REINSURANCE PRICING: PRACTICAL ISSUES & CONSIDERATIONS

8th September 2006

2006 GIRO “Reinsurance Matters!” Working Party

Mark Flower (Chairman); Ian Cook; Craig Divitt;
Visesh Gosrani; Andrew Gray; Gillian James;
Gurpreet Johal; Mark Julian; Lawrence Lee;
David Maneval; Roger Massey

Abstract

The purpose of this paper is to offer a collection of practical thoughts and considerations relevant to the pricing of reinsurance contracts on (predominantly) liability classes of business in the London Market. The paper is aimed at the practitioner and as such assumes a working knowledge of experience and exposure rating. We aim to supplement existing texts by pointing to some of the best relevant papers and then discussing life in the world of tailored reinsurance contracts, poor information, complex heterogeneous exposures and limited evaluation time.

Don't panic!

This is a fairly long paper but you don't have to read it cover-to-cover as it isn't one single narrative. We hope you will still be able to gain something useful by skimming though quickly and dwelling on the areas that particularly interest you. The paper is also largely non-technical.

Thanks!

The members of the working party would like to express their gratitude to Martin Burke, Carl Overy and Stuart Wrenn for providing insightful, helpful and polite feedback on review of our draft paper. This assistance was much appreciated – thank you.

All errors, omissions and potential controversy however remain the sole responsibility of the authors.

CONTENTS

0) EXECUTIVE OVERVIEW	4
1) INTRODUCTION AND PRELIMINARIES	5
2) BASIC PRICING TECHNIQUES	8
a) Experience rating (Burning cost)	
b) Exposure rating (ILFs etc)	
c) Hybrids (Frequency / severity)	
3) SOME THOUGHTS ON PERFORMING THE BASICS	10
a) When using stochastic models...	
b) Curve fitting	
c) Statistical distributions: which ones when and why?	
d) Combining experience / exposure / market / expiring rate	
e) LAE Treatment (Pro rata vs. UNL/costs inclusive.)	
f) Underlying profitability	
g) Risk loads	
h) Rate Increase Analysis Considerations	
i) ECO / XPL	
j) Cession treaties need pricing too!	
k) Don't ignore the words	
l) Slip rates of exchange	
4) DATA ISSUES	29
a) No data – e.g. high excess layers	
b) When “As-If” becomes “If Only”...	
c) Tips for flushing out data issues	
d) Other practical miscellany	
5) PRICING “BELLS AND WHISTLES”	34
a) Limits on Cover: AADs, Corridors, Mandatory Co RI, Aggregate Caps	
b) Retrospective Premium Adjustments: Swing Rates, RPs, APs, PCs, NCBs	
c) Franchises and Reverse Franchises	
d) Top & Drops	
e) Multi-year vs annual	
f) Indexation clauses	

Continues....

6) UNDERWRITERS' RULES OF THUMB – HOW TO TEST THEM	49
a) 1x1 as rate on 1x0, 1x2 as rate on 1x1, etc	
b) Discounts for aggregate caps or limited / paid reinstatements	
c) Minimum Rates on Line	
7) COMMON PITFALLS	54
a) Ignoring drift in underlying limit / deductible profile	
b) Using ILFs without considering currency, trend, LAE, etc	
c) Pricing Retro treaties on ventilated programs without “looking through”	
d) Trend issues	
e) Extrapolating past years of a fast growing account	
f) Ignoring “Section B” (e.g. clash or other “incidental” exposure)	
g) Changes in underlying coverage e.g. Claims Made vs Occurrence	
8) ACCUMULATIONS AND CORRELATIONS	59
a) Pricing Clash	
b) Correlations for the Real World	
c) Systemic Losses	
9) CREDIT RISK ISSUES	64
a) Pricing considerations	
b) Modelling credit risk	
c) Other considerations	
10) SO WHERE DO WE GO FROM HERE?	69

APPENDIX

A LITTLE MORE ON CREDIT RISKS...	70
----------------------------------	----

0) EXECUTIVE OVERVIEW

- This paper is not a basic instruction manual, it is more of an ordered collection of our practical thoughts and suggestions that we hope will be of interest to the London Market reinsurance actuary. It is more a “reference book of tips and suggestions” than “this is how you should do it”. It also has a liability focus, although many of our points apply equally to property.
- Our two over-arching beliefs are set out in section 1: “intelligent execution” and that old chestnut “communication”. *Intelligent Execution* means remembering the context of your information and environment and focusing most of your time on getting a solid grasp of the key dynamics involved, rather than spending too much time worrying about theoretical purism around the edges. *Communication* is as much about convincing people of what you *don't* know and what you *can't* infer, as it is about what you do and can.
- The bulk of the paper then gathers our combined thoughts and tips for “pricing success in the real world” into eight sections:
 - §2 refers to some good existing papers explaining the basic pricing methodologies;
 - §3 explores 12 specific aspects of performing the basics, from curve fitting and stochastic models, through treatment of LAE and pricing of ECO/XPL exposures to reading the submission;
 - §4 deals with common data issues you will be faced with, such as considerations when modifying the experience to create an “as-if” position;
 - §5 lays out seven categories of common reinsurance contract design features that influence value, and discusses how you might factor these into your pricing analyses;
 - §6 covers the continued use by underwriters of “rules of thumb”, such as minimum rates on line or apparently made-up discounts for last minute adjustments to the coverage. We don't try to give a long list, but we do encourage you to challenge these and illustrate how you might go about this;
 - §7 reveals a large number of pitfalls when pricing reinsurance and suggests ways to avoid or mitigate them;
 - §8 is about accumulations and correlations. They exist, they affect the value of reinsurance contracts and as such they should be considered – even if doing so is challenging and you don't have the data to do it rigorously. Again we offer some ideas for the real world;
 - §9 contains a brief discussion of credit risk issues, and we also suggest a possible way to quantify credit risk using CDS prices. We are only scratching the surface here, and we believe there is a need for more research and practical papers in this arena.
- Our last chapter, §10, looks to the near future. We suggest two areas where we perceive particular value in additional practical research and development, being accumulations/correlations and the credit risks just noted. We also note two areas where we believe the Faculty and Institute of Actuaries might lend its weight to advance matters, firstly in the arena of interactive educational resources focussed on practical issues, and secondly in the development of benchmarks from market data.

1) INTRODUCTION AND PRELIMINARIES

This paper is not intended as a Complete Beginner's Guide to Reinsurance. We are targeting an "intermediate" audience, actuaries already familiar with the key concepts and basic methodologies in common use, although that does not mean that others won't also find some of this of use or interest. There are plenty of good papers out there covering the fundamentals and we will reference these rather than replicate them. Incidentally, there are also some good books on the subject of Reinsurance, such as "Reinsurance in Practice" by Robert and Steven Kiln.

Before we got stuck into writing this paper, we first trawled the e-libraries and dusted off our own reading piles to prepare a "Reinsurance Actuary's Reading Guide". This is a list of 42 actuarial papers relating to reinsurance (with a predominantly liability flavour) and we have compiled a very brief overview of each of the 42. The primary purpose of this exercise was to take stock of the great resource that already exists to ensure we didn't waste time reinventing the proverbial wheel. However, we also hope that this reading guide will lead to a reinsurance addendum to the official Faculty & Institute Reading Guide. Our guide is imaginatively entitled "Review of papers relevant to non-life reinsurance (with a liability focus)" and is available from any one of the working party or via the 2006 GIRO Conference Papers area of the Institute of Actuaries website¹.

Our objective in *this* paper is to build upon all the existing actuarial papers by offering our **practical** thoughts and considerations on a range of related topics as listed in the contents page. Our paper is not structured as a "how to do something" manual and as such may not lend itself to being read cover to cover. However, we have arranged our thoughts into categories by way of a reference guide. We hope it will be of interest to the practitioner when dealing with some of these topics – in particular it should give someone relatively new to the pricing of reinsurance a good overview of many of the issues and considerations you will have to deal with.

We have focussed our comments primarily in the context of long tail liability reinsurance, although some of the issues are equally applicable to the short tail arena. Note in particular that we have not considered the property cat models in any way.

As a general rule, topics such as alternative risk transfer (ART) and non-traditional reinsurance contracts, accounting issues, the re/insurance cycle, the pricing/reserving loop, reinsurance as vicarious capital, capital allocations and risk-adjusted return are all largely beyond the scope of this paper – we had to draw the line somewhere, even though many of these topics are worthy of review.

It is important to recognise that the reinsurance pricing world is far from perfect. Information is invariably less good than you would like, and as such this is not an arena in which the purist theoretician will feel especially comfortable. There are many good papers of elegant theory that unfortunately, being based on neat assumptions such as homogeneity and stability can sometimes be hard to successfully harness in our highly imperfect reinsurance world. Our paper is not of that ilk.

However, all of this does not mean that the actuary cannot add value! The application of actuarial techniques will invariably involve taking short cuts and making approximations and simplifications in order to get the job done, but we are still in the business of interpreting data in order to offer insight, perspective and guidance based on methodologies, techniques and assumptions that are fit for purpose.

¹ http://www.actuaries.org.uk/Display_Page.cgi?url=/library/proceedings/gen_ins/gic_index.html

We believe that two keys to success are Intelligent Execution and Communication:

Intelligent Execution

Whatever you are doing, remember the context in which you are doing it, the purpose for which you are doing it, and the needs of the end user (usually an underwriter). Make sure you've got the basics right, you understand all the key dynamics and you respect the limits of your information and analysis.

Communication

As with all fields of actuarial work, effective communication is of paramount importance. Obviously that means explaining your analysis, assumptions, sensitivities and weaknesses and your viewpoint in a readily digestible manner. Only when the underwriter (or CFO or whomever) fully understands and accepts what you are saying should she take any notice of it.

But we think it's actually a bit more subtle than that with reinsurance, particularly in the London Market.

- There are still those who are implicitly sceptical of actuarial input: for these people the style in which you communicate is often as important as the message you are communicating. This is nothing new.
- However, it seems there are also an increasing body of people who are too accepting of what the actuary tells them. This might be for honourable reasons (they trust you unquestionably) or for more dubious reasons (they want to hide behind you as a reason for their actions, or they simply don't plan on taking any notice of what you say anyway so why bother to argue) but either way it isn't a healthy situation.

Read the signals to improve your communication. If you are not being challenged on your assumptions and the underwriter is putting undue faith in the precision of your (lets face it, sometimes wobbly) conclusions this will come back to haunt you.

It is important that your customers understand the limitations of your work product, and it is also important that you and they either agree on all the key assumptions or you both know where and why you disagree, and what that difference of opinion means in terms of the end result.

There was a pertinent article printed in the Economist magazine in January 2006. It was framed in a Pensions context but is nevertheless applicable to reinsurance actuaries and serves to illustrate the point:

“AMONG the many jokes about actuaries, one cruelly hits the mark. An actuary and a farmer are looking at two fields of sheep. The farmer asks the actuary how many sheep he thinks there are: “1,007”, is the quick and confident reply. The astounded farmer asks how the actuary reached that number. “Easy, there are seven sheep in that field and about 1,000 in the other.”

False precision and reckless approximation have defined the actuarial profession's role in the crisis that has enveloped corporate pensions on both sides of the Atlantic. Although actuaries have not been the only cause—companies, trustee boards, governments and accounting rules have all played their part—they have been surprisingly hapless at their main task: forecasting funds' future liabilities and assessing how many assets will be required to meet them.”

In this paper we will walk quickly through the various topics listed in the contents page, referring to pertinent existing texts and offering our observations in the context of intelligent execution and communication. You may well find some of this paper to be obvious, but we hope that you will also find other parts to be interesting, thought provoking and applicable in the real and imperfect world of reinsurance pricing.

2) BASIC PRICING TECHNIQUES

- a) Experience rating (Burning cost)
- b) Exposure rating (ILFs etc)
- c) Frequency / severity based rating

By “Experience Rating” we mean the practice of taking contract-specific historical loss and exposure data, developing it to ultimate, trending and “on-levelling” it to better reflect prevailing exposures and using this as a basis for selecting your expected losses and hence required premium rate for a reinsurance contract.

By “Exposure Rating” we mean the practice of using the reinsurance exposures together with general industry data about loss ratios and severity patterns as a basis for selecting your expected losses and hence required premium rate. Typically the exposures will be quantified in a policy limit / attachment profile together with some details of the underlying policies, i.e. class and sub-class of business etc. The industry data will usually take the form of ILF curves (Increased Limit Factors) or something similar to allocate losses between layers of exposure, together with a loss ratio expectation for the underlying book as a whole.

By “Frequency Severity Rating” we mean the practice of developing a stochastic model of the risk where you simulate the potential loss experience based on estimates of the number and size of reinsured losses together with the reinsurance treaty terms. This gives you a distribution for the recoveries and thus leads to a price. Such a model can be very simple (e.g. with @Risk in a simple spreadsheet) or very complex (e.g. using DFA software). Stochastic modelling sits somewhere between the experience and exposure methods in that the key assumptions might be based either on the contract’s experience, on more general industry data or on a combination of the two, depending on what you have to work with.

We do not propose to give an explanation on how to perform these pricing analyses as this is already well documented. For example, a good foundation can be had by reviewing the following:

- 1995 GISG Pricing In The London Market (Sanders et al)
Offers the novice a good overview of the basic techniques – burning cost, frequency / severity analyses, exposure rating with ILFs and using credibility theory in practice. It also offers a good guide to Profit Testing, which can be useful when analysing larger contracts.
- 1997 A Simulation Approach in Reinsurance Pricing (Papush et al)
This could be renamed “How to start building and using frequency / severity rating models”
- 1998 GISG Reinsurance Pricing (Chandaria et al)
This offers comment on some of the technical issues encountered when pricing reinsurance, including trend, LDFs, exposure data, stochastic models and loss sensitive contract terms.
- 1999 GISG Reinsurance Pricing (Chandaria et al)
Complementing the 1998 paper, this highlights a number of pitfalls, including data thresholds, IBNE, underlying policy limits, extreme losses and using the wrong aggregate loss distribution.
- 2004 GISG Reinsurance Pricing – Rating Long Tail XoL (Cockroft et al)
This paper offers a second perspective on many of the technical issues in the above papers.

Incidentally, there are also many good papers on more specific topics, such as:

- 1977 Miccolis paper On The Theory of Increased Limits and Excess of Loss Pricing
Don't dismiss this one simply because it's so old... it offers a good explanation of how ILFs are derived and used.
- 1981, 1989, 1994 and 2000 CAS papers on Trend
Respectively, these years saw CAS papers published covering the impact of trend on severity curves, layer losses, how to choose a trend factor and loss trend considerations when experience rating.

A quick search of the public e-libraries found on the websites of both the Faculty and Institute of Actuaries and the Casualty Actuarial Society is all you need to do to find these and many other papers, or have a look at our “Reading Guide” as mentioned in section 1.

For the rest of this paper we assume a working familiarity with experience, exposure and frequency severity rating techniques. Our approach from hereon in is simply to whiz through a host of related topics, offering our collective thoughts and comments that we hope are of practical use and/or interest to the actuarial reinsurance practitioner.

- For example, you might have used a LogNormal when you could have chosen a Pareto with similar first and second moments and apparent goodness of fit to your imperfect data. Swapping from one to the other might make little difference to aggregate expected losses, but it could have dramatic implications for your assessment of recoveries on high excess of loss layers, or on your capital allocations, or on your views about the depth of aggregate reinsurance coverage.
- One interesting paper on this topic is “Estimating the Parameter Risk of a Loss Ratio Distribution” presented by Charles Van Kampen to the CAS in 2003. This considers parameter risk due to selecting one from a number of possible curve fits to a dataset (but note that your modelling and parameter risk could be wider reaching than this).
- We recommend performing sensitivity tests on key outputs relative to your model design, assumptions and correlations, and handling all your output with due care. Don’t spend all your time tinkering with the “trees” when building the model, think about the “wood” too.
- If the data upon which you parameterise your model has no large losses above a point such that you are nervous about your severity fit, consider blending your stochastic model with more deterministic approaches. For example, use the model to generate losses in the first layer of coverage only (or part thereof) but then apply ILFs to this output in order to assess the higher layers.
- Correlations can have a dramatic effect of the results generated by a stochastic model. If you ignore correlations you are likely to significantly underestimate the tails in your downsides, which is usually A Bad Thing. However, correlations are just about the hardest part of a stochastic model to do well because there is so little data to guide you. Correlations warrant a closer look so please take a look at Chapter 8 (Accumulations and Correlations) in this paper for a little more on this subject...

Stochastic modelling seems to us like an area that could benefit from a new “practitioners’ guide” to update Dmitri Papush’s 1997 paper to contemporary best practice. We suggest this might be a valuable project for a GIRO working party?

b – Curve fitting

Simulations based upon an empirical distribution can be sufficient under certain conditions (i.e. where you have a lot of data) but this isn’t usually the case when dealing with reinsurance contracts. We will normally be in the business of fitting and selecting statistical distributions.

There are any number of probability and statistics text books describing techniques for fitting curves to data, and this is also covered by the actuarial examinations:

- Using QQ-plots and mean excess functions to suggest curve families;
- Parameter fitting techniques such as moment matching;
- Goodness of fit tests such as Chi-squared;
- Dealing with trend, truncation etc.

In practice these days the curve fitting process is all streamlined and the mathematics can be handled rapidly by curve fitting software. This is obviously a great help, but we implore you:

Don't get lazy!

The process of fitting a distribution should not be fully automated...

- With a small dataset and limited market information a variety of distributions might provide “a good fit” but actuarial judgment is of paramount importance – a seemingly good fit to imperfect data is **not** necessarily a good model to use.
- As good as your fitting software might be, it is not infallible and often the parameterisation routine can result in fitted curves of the “could do better” variety. For example, one particularly common package fits a two-parameter pareto by fixing one parameter at the smallest observed loss and then solving for the best fitting second parameter. Quite often you could get a much better fit by using a different first parameter, but the software won't do this.
- Recognise that with small datasets your fitted parameters are likely to be biased. For example, your data could well be devoid of any very large losses and your fitted mean and variance could be understated.
- Before fitting the curves, we usually go through a process of data cleaning and adjusting “on level” (i.e. adjusting to allow for differences in the forthcoming exposure period relative to the past years, such as for claims inflation and changes in the volume and profile of exposures.) This process in itself has a significant impact on the curves you will select and thus on your results and conclusions, and as such warrants due care and attention.
- There's really no substitute for using your eyes and engaging your brain:
 - How good does the fit look on a quantile-quantile plot? Look at the fit over the full range, but also focus on specific areas that matter such as the tail – especially if you are fitting a loss severity distribution to model excess of loss recoveries;
 - Truncation of loss amounts by policy limits will lead to “clustering” in your loss data at certain key points of the distribution, for example at \$1m. It is important to capture this effect in your model, but an automated curve fit will usually not be able to handle this;
 - Don't over-fit your model to your data. For example, when fitting loss severities you should expect to have a relatively poor fit at the upper end of the PDF where the empirical data is patchy / lumpy but the parametric curve is smooth. If you are going to over-fit you may as well just use the empirical distribution;
 - Compare the fitted tails with the extremes from your data. For example, if you are fitting a loss frequency distribution using 10 data points (i.e. 10 annual loss frequencies adjusted on level for the coming year's exposures) how do the smallest and largest data points compare with your fitted 5th and 95th percentiles? They don't have to be the same, but they should make sense based on what you know;
 - When fitting a loss severity distribution for a class of business, does it seem reasonably consistent with the ILFs your underwriter uses for that class? If you don't know the maths you can readily test it using a simple @Risk model – simulate a loss

capped at a limit such as \$1m and note the Limited Expected Value (LEV²). Repeat with, say, a \$5m limit. Compare the ratio of LEVs (5 over 1) with the ILF for \$5m versus the ILF for \$1m. If your LEV ratio is a lot higher than the ILF you have chosen a relatively severe curve, and vice versa;

- Use the Chi-squared and other such numerical tests with care: a visual inspection of the fit will tell you more than reducing the comparison to a single numerical measure.
- Above all, think about what you are doing – does it make sense?

c – Statistical distributions: Which ones when?

Commonly Used Frequency Distributions

- **Poisson:** The Poisson distribution is often used to model loss frequencies. It is very easy to parameterise and it is also additive (i.e. $\text{Poisson}(a) + \text{Poisson}(b) = \text{Poisson}(a+b)$) which can sometimes be helpful in simplifying your models.

However, the Poisson involves the implicit assumption that all of your insured risks are independent. It is also a “memory-less” distribution, so for example having a loss event in the first half of the exposure period tells you nothing about the likelihood of a subsequent event in the second half. And as we all know, the Poisson has a variance that is equal to the mean.

There are some circumstances where these assumptions hold in entirely fortuitous classes, such as fire insurance on a property book (maybe...). However, and especially so in liability insurance, these assumptions are often invalid in which case the Negative Binomial distribution should be a better choice as it does away with these assumptions.

- **Negative Binomial:** The NB is a generalisation of the Poisson, and typically has a longer tail and greater variance.

The additional flexibility and the process of explicitly thinking about the variance as well as the mean are both helpful. When modelling working layer mono-line liability reinsurance, using a NB with a variance of around three times the mean would not be unusual.

Quite a lot has been published in the actuarial arena on the Negative Binomial distribution, much of it dating back to the early 1960's. If you are interested in a quick history lesson take a look at Simon LeRoy's short papers appropriately entitled “The Negative Binomial and Poisson Distributions Compared” (1960) and “An Introduction to the Negative Binomial Distribution And It's Applications” (1962).

It is also handy to remember that the NB distribution is actually a “Poisson with a Gamma mean”, i.e. a Poisson distribution with a random parameter that is Gamma distributed. This can sometimes be used to your advantage, for example if you are simulating with a package that constrains you unduly when setting the NB parameters.

² LEV – The expected value of the capped loss distribution. For example, the LEV at £1m would be the expected value of $\text{MIN}\{x, £1m\}$

- There are other discrete distributions you may well use, such as an over-diversified Poisson, and you also shouldn't feel constrained to using the discrete distributions readily to hand:
 - It doesn't take a lot of effort to "stratify" the density function of a continuous distribution and thereby create a discrete version of it if you think it a better curve to use (e.g. a discretized Gamma);
 - You may want to manually adjust any chosen frequency distribution to allow for any perceived exposure features that are not already captured in your data and fit.
- The main thing is to think about what it is you are modelling and whether or not your fitted distribution seems to make sense across the range of possible outcomes, not just at the mean.

Commonly Used Severity Distributions

Perhaps the most common severity distributions currently in practical use are as follows. (Please refer to your statistics text books or the internet if you'd like associated formulae.)

- **LogNormal:** This distribution is in widespread use for modelling ground-up loss data (i.e. all losses, not only the large losses). It is not only the right "shape" to fit the loss experience of many insurance classes, it is also particularly easy to work with in a spreadsheet as the various formulae you'll want to use for probabilities, limited expected values and so on are very manageable.
- **Pareto:** Again this family of distributions is in widespread use, and also comes in a variety of forms from an easy to use single-parameter version to a more flexible but cumbersome five-parameter version. The Pareto has a thicker tail than the lognormal and is usually fitted to "large loss data", i.e. losses above a threshold. Two forms in particular common use are the "single parameter Pareto" (because it's very easy) and the five-parameter Pareto (which was used frequently by ISO in their circulars based on industry data, although latterly they seem to use the Mixed Exponential).

For a good overview of working with the single parameter Pareto we recommend you read Philbrick's 1985 paper "A practical guide to the single parameter Pareto distribution".

- **Mixed Exponential:** This is simply a blend (or a weighted average, if you prefer) of a number of exponential distributions. The ISO use it nowadays for their ILF development, as they believe it gives an easier and better fit to insurance data than does the generalised Pareto. You can blend as few or as many as you like – ISO typically blend between 5 and 8 curves.

The 1999 Keatinge paper "Modelling losses with the mixed exponential distribution" is a good place to find out some more about this family.

- **Beta:** This is a medium-tailed distribution, and one of the few distributions for which you can directly calculate the value between upper and lower truncation points. This makes it easy to work with for reinsurance, as it is possible to calculate both the mean and standard deviation within a reinsurance layer for any one loss.

The Beta distribution is used by RMS to model the severity of a specific event; as each event has a predetermined annual rate you can then calculate the mean and standard deviation of the aggregate loss cost in a reinsurance layer, albeit for a contract with unlimited, free

reinstatements. In practice most actuaries will need to assess reinsurance with features that limit the aggregate recovery and / or retrospectively adjust the premium based on loss experience, so will use this distribution within a simulation model rather than doing the maths.

- There are clearly other distributions used by some practitioners, including the Gamma, Burr and Weibull.

We feel the key to success here is how appropriate your fitted distribution is for the purpose you wish to use it for, rather than how tightly it might fit your data. In many circumstances your raw data will not adequately describe the true forthcoming severity profile for a number of reasons, some common ones being:

- Not enough (or too many) very large losses in the experience leading to a misleading picture of the distribution tail. What does your curve suggest about the return period of a loss bigger than, say, £1m? £5m? £10m? Does this seem reasonable? Is your underwriter convinced?;
- Changes in the underlying exposure profile, such as an increase in average policy limits, or a tort reform, or a shift in the proportion of, say, obstetricians and gynaecologists in a medical malpractice book;

As such it would be relatively unusual to use an empirical distribution, but for the same reasons you should consider seeking a distribution that does not fit your data in a specific way. For example:

- If you think your data is too light in the tail, use a thicker tailed distribution and be aware that this may well lead to a higher average cost per claim versus your data;
- Consider blending distributions and / or adding synthetic losses from elsewhere, e.g. simply using your judgement or by sampling from a set of scenarios such as the Lloyd's RDS; This is straightforward to implement in a spreadsheet environment.

Finally, whilst talking severities, we should mention EVT (extreme value theory). This is based upon fitting a generalised pareto distribution to extrapolate the tail of a distribution based upon data for very large losses. Whilst much interesting research has been done in this area, it is important to remember the imprecision of fitting any model to sparse data. EVT was en-vogue a few years back so there are a number of recent articles on the subject, but for the actuarial reader in particular we suggest you take a look at Sanders et al's paper from 2002 entitled "The Management of Losses Arising from Extreme Events". (Quite a lengthy paper including a lot of additional background material but Section 3 describes the application of EVT to insurance.)

However, unlike the banking world, EVT does not seem to have shaken up a lot of reaction in the reinsurance arena as yet. One possible reason for this is that we have been using heavy tailed distributions for a long time and EVT is simply another twist in a familiar path whereas until recently much of the banking community used only normal distributions in their modelling.

d – Combining experience / exposure / market / expiring rate

Okay, so you’ve done all your analysis, allowed for your internal expenses and profit targets etc and come up with a set of premium rates. You have your experience and exposure rates, your broker or underwriter tells you what the market rate is likely to be, and you also know what the rate was last year. You might also have some additional comparison rates base on similar risks. How do you fit them all together?

A few papers have been published seeking a technically valid credibility measure to weight experience and exposure rates, most recently “Bayesian Credibility for XoL Reinsurance Rating” by Cockroft (2004). However, as good as such approaches might be for “well behaved” business, they are less easy to apply rigorously in the “badly behaved” world of reinsurance.

Blending experience with exposure

- If you want to blend your reinsurance experience and exposure rates using a credibility factor we suggest you apply something simple and pragmatic, accepting that it is going to be imperfect.
- One such approach (analogous with the Bornhuetter-Ferguson reserving technique) could be based upon your loss development factors for the class, for example:

Experience rate based on, say, 5 years of data for which we have:

Year	LDF	% Reported	On-Level Exposures	Credibility Factor
2001	1.10	90.9%	80	13.7%
2002	1.25	80.0%	90	13.6%
2003	1.75	57.1%	110	11.9%
2004	2.50	40.0%	120	9.1%
2005	5.00	20.0%	130	4.9%
			530	53.1%

Given that 2001 is 90% developed then it is in some way “90% credible”, but based on its exposure volume of 80 part of the overall 530 it contributes 15.1% of the experience based estimate – so 90% credibility x 15.1% of the estimate = 13.7% credibility factor. Overall your experience estimate is “53.1% credible” by summation. This calculation is quick, easy and intuitively pleasing, although no doubt the purist will poke holes through its technical validity.

- Of course, just because your experience rate might be “53% credible” doesn’t mean your exposure rate is necessarily 47% credible. Your exposure rate might be far better or worse than that as the two measures are not intrinsically linked in this way. We are not suggesting you use our “back of the envelope” credibility factor, we simply encourage you to think about how credible each estimate is and seek a methodical way of allowing for this.

Market rate, expiring rate, peer comparisons

The Market Rate, Expiring Rate (last year's premium) and Peer Comparison Rates (what you charged on the most recent similar contracts bound) give you a context for your technical rates. (We use the term "technical rates" here as a collective noun for your experience rates and exposure rates, although you will no doubt encounter varying uses of this phrase whilst working in reinsurance!)

- The market rate can be quite different to your technical rates, for a number of reasons:
 - The lead underwriter(s) might have made a different assessment of the risk based on their own knowledge and experience. Remember, there are a lot of subjectivities involved so you might simply be out of line with others in your thinking;
 - This might be cycle related: in a hard market, the underwriter is taking advantage and writing for super-profits, or in a soft market the broker is attempting to hold the underwriter's toes to the fire;
 - Sometimes there will be significant terms and conditions or other more qualitative features of the risk that are factored in to the market rate (up or down), but which do not feature in your technical calculations;
 - There may be commercial considerations at play, such as writing one contract at an expected underwriting loss (or minimal profit) in order to access other contracts that generate sufficient expected profits that it makes good sense to write the full set. This can happen quite frequently on a reinsurance program where one broker is placing several layers of reinsurance for the same client at the same time and the broker will only allow a market to participate on the "juicy" layers if they also take a share in the more aggressively priced layers.
- Given the cyclical nature of market rates it is worthwhile tracking your technical rates as a consistent benchmark against which to monitor. It is hard to infer anything from individual risks deviating from benchmark, but observing an entire portfolio moving up or down can be very instructive, not to mention useful for your reserving actuaries!


Which ones when?

- Remember, all of this analysis is to some degree subjective and based on imperfect information, so none of the various rates will be "right" in the technical sense.
- Furthermore, each party to the process has a different bias and as such will come up with a range of different prices for any given case:
 - The buyer wants a cheap price and believes sincerely and passionately in the quality of their risk management;
 - The sellers (reinsurers) want a profitable rate (translation: expensive price) and have a cynical eye for potential downsides. Note that with syndicated placements (several reinsurers each taking a share of one reinsurance risk) this will generate several different prices, often associated with several different requirements regarding the contract terms and conditions; (i.e. they will negotiate contract terms as well as prices);

- The broker attempts to negotiate a settlement between these parties with an emphasis on securing the best deal possible for his client, the buyer. “Best deal” generally means “cheapest price at which he can complete the placement whilst also satisfying the client’s requirements regarding depth of coverage and quality of the reinsurance panel.”
- Perhaps one intelligent execution of all these various rates is to consider all the rates side by side rather than try to meld them into one “right answer”. From your analysis of the risk you should have a feel for how credible each of your various rates are so there is no particular need to blend them into one, and doing so can actually hide information such as the differential between the experience and exposure rates.

Compiling a simple exhibit showing your various rates side by side can give you and your underwriter a much better feel for the appropriateness of the market rate than simply lumping it all into one final blended “answer”. This might be very simple, for example:

Actual Burning Rate	: 18.3%	<i>Incurred losses / Subject premium...</i>
As-If Burning Rate	: 14.2%	<i>...adjusted for classes no longer written...</i>
Experience Rate	: 21.6%	<i>...then trended and developed.</i>
Exposure Rate A	: 24.1%	<i>Based on their ILFs</i>
Exposure Rate B	: 25.4%	<i>Based on our own ILFs for this class</i>
Expiring Rate	: 23.5%	<i>Price charged last year for same cover</i>
Proposed Rate	: 22.5%	<i>Broker says it will be placed at this price</i>



However, beware – there is always a risk when presenting underwriters with such exhibits that if there is a wide range in the apparently reasonable answers you (a) give them scope to charge whatever they want and attribute this to your advice, and (b) can expect the underwriter to react cynically if you don’t give a good explanation as to why all of your numbers vary as such.

e – LAE Treatment (Pro rata vs. UNL/costs inclusive.)

Not all excess of loss reinsurance treaties handle LAE (loss adjustment expenses) the same way, which has direct and significant implications for the likely recoveries and hence the value / price of the contract. The two common approaches are “Pro Rata In Addition” and “Costs Inclusive” (note that Costs Inclusive is sometimes referred to as “UNL” or “Ultimate Net Loss”):

- Pro rata in addition: Indemnity is first allocated to layers according to the limits and attachment points, and then LAE are allocated in proportion to the indemnity.
- Costs inclusive: LAE is included within the definition of loss (i.e. loss = indemnity plus LAE) and the limit and attachment point are applied to this sum total.

For example, consider a \$1.4m insured loss comprising \$900k of indemnity payments and \$500k of LAE. Suppose the insurer buys two reinsurance layers: \$500k xs \$500k and \$4m xs \$1m. The two allocation bases work out as follows:

		<u>Gross Loss</u>	<u>4 x 1</u>	<u>500 x 500</u>	<u>Net Loss</u>
Pro Rata In Addition:	Indemnity	900,000	0	400,000	500,000
	Expense	500,000	0	222,222	277,778
	Total	1,400,000	0	622,222	777,778
Costs Inclusive:	Total	1,400,000	400,000	500,000	500,000

This simple example illustrates that:

- With pro-rata in addition a reinsurance layer can pay significantly more than the “limit” for any one loss, but it tends to “squash the allocation downwards” relative to the UNL basis;
- If you buy reinsurance on a Costs Inclusive basis, you need to buy limits well in excess of your maximum policy limit in order to reduce the chances of running out of vertical cover;
- If you buy reinsurance on a Pro Rata In Addition basis, you are exposed to retaining significant LAE on cases where the indemnity finally settles for a relatively small amount.
- The expense allocation basis is a fundamental part of defining the contract, almost as important as the limit and attachment. As such you need to understand what you are dealing with and allow for this in your pricing analysis.

One common related issue is the application of ILFs when exposure rating – are your ILFs “indemnity only” based or “UNL” based? How does the reinsurance treaty work? Taking ILFs developed on a UNL basis and applying them to a limit profile to exposure rate a pro-rata-in-addition basis treaty, for example, will tend to allocate relatively more exposure to higher layers.

Here is a simplistic illustration:

- Suppose you have only one insurance policy which provides a policy limit of \$5m indemnity together with unlimited associated LAE. Based on an (indemnity-only) loss severity curve for the class of business, suppose further that the expected indemnity payments fall 60% in the first \$1m and 40% in the layer \$4x1m. If we also expect 50 cents of LAE for each dollar of indemnity and we allocate the LAE “pro rata in addition” to the indemnity we have the following table of expected indemnity and LAE by layer and thus some implied ILFs:

	Pro Rata In Addition Basis			
	Indemnity	LAE	ILF	
5x5	0	0	n/a	
4x1	40	20	1.667	
1x0	60	30	1.000	
Total	100	50		

Applying these ILFs to the \$5m policy limit to exposure rate a 4x1 layer (again with LAE pro-rata in addition) would return the “correct” rate of 40% of premium:

$$(1.667 - 1.000) / 1.667 = 0.400 = 40\% \text{ rate}$$

However, contrast this with how the losses and expenses might alternatively be allocated on a UNL basis, and thus how ILFs prepared using UNL data might look. With UNL we consider each loss in its entirety (i.e. the indemnity and LAE added together) and thus have a “longer” severity curve stretching into the 5x5 band. Using this “UNL basis” severity curve to allocate the total costs we might have a “UNL Basis” allocation and ILFs that look something like this:

	Pro Rata In Addition Basis			UNL Basis	
	Indemnity	LAE	ILF	L+LAE	ILF
5x5	0	0	n/a	10	2.143
4x1	40	20	1.667	70	2.000
1x0	60	30	1.000	70	1.000
Total	100	50		150	

Using these UNL-basis ILFs to exposure rate the same, “4x1 with LAE in addition” policy we would reach the very different exposure rate of 50%:

$$(2.000 - 1.000) / 2.000 = 0.500 = 50\% \text{ rate}$$

Conversely, using “costs inclusive ILFs” to price a “costs in addition” treaty will also lead to a mismatch and thus a potentially significant mis-price.

The moral is simple: make sure you know the LAE basis of your reinsurance contract, but also make sure you know the LAE basis of your ILFs!

f – Underlying profitability

This is a very simple but also very important point that is all too often ignored, often because the broker won't automatically think to provide the information you would need to assess this. The onus is on you to remember to look for it.

If you are rating an excess of loss contract, the premium will usually be expressed as a function of the underlying subject premium, for example "21% of Gross Net Written Premium Income". As well as verifying whether you believe 21% is a fair rate for the layer, you should also validate the underlying profitability. Suppose 21% is a fair allocation, but the underlying book is running to a 95% gross loss ratio and you are targeting a 75% loss ratio. Clearly a 21% premium rate is unlikely to do this for you, and you should therefore seek something closer to $21\% \times 95/75 = 26.6\%$.

Just because the reinsurance rate might have increased marginally year on year, doesn't mean your loss ratio expectations are improving. Loss trend, limit / attachment profiles and also the profitability of the underlying business all have an important impact.

Taking this one step further, if you have the time and the information to do so, it can sometimes be instructive to have a quick look at the overall profitability and rate levels of your cedant, for example in the case of a US stock company you might review Schedule P of their Statutory annual reports. This should help you understand your client a little better, it may help you in the context of exposure level adjustments and it could also give you a different perspective on why they might be buying the reinsurance you are selling. (If your cedant is a stable, profitable giant can you think of good reasons why they might be interested in buying your reinsurance other than opportunistic, i.e. they think your pricing is cheap??)

g – Risk loads

Much has been written on risk loads for reinsurers, and there are a number of approaches in practical use. These range from simple measures based on standard deviations (or variance or CoV) to more complex DFA based measures such as a target return on risk based capital.

For a good introduction, we would encourage the reader to review the paper "Risk Loads for Insurers" written by Feldblum in 1990, together with the subsequent discussion papers issued by Philbrick, Bault and Meyers. Also well worth a look are the following papers:

- Implementation of PH-Transforms In Ratemaking: Wang, 1997 (and discussions);
- Risk Load and the Default Rate of Surplus: Mango, 1999;
- Capital Consumption: An alternative method for pricing reinsurance, also by Mango, 2003.

A detailed analysis of the alternative approaches is beyond our scope here, but we would simply offer the following comments:

- Whatever you are doing, keep it in context. What metrics are being used by your management to allocate capital, measure underwriting returns and remunerate the underwriting team? It is very important to align your risk loading methodologies with these metrics.

- Don't over-complicate matters. There will usually be significant uncertainty around the expected recoveries from a reinsurance layer, so a complex risk load that takes a lot more effort to calculate and communicate than a simple approach, but which offers only a small increase in "theoretical robustness", may not be the best one to use. It could also undermine your credibility with the underwriter if they struggle to understand and accept your approach.
- We understand there is a now bit of a move by some markets towards more sophisticated measures, although this is not universal.
- Beware the shape of your curves. Two risks with the same mean and standard deviation but different shape PDF probably warrant different risk loads (i.e. a straightforward % SD load is not robust). But see the immediately preceding point before getting too hung up on this one!
- Different players (the cedant and each reinsurer) have different diversification capabilities and thus should use different risk loads. Treat your risk loadings as a personal thing and don't just copy your neighbour. (Simon Pollack and Dean Dwonczyk presented a workshop at GIRO 2004 with a more detailed discussion of this concept – their papers are on the IoA website.)
- Consider loading your expected loss cost to reflect the asymmetry of the data in the relationship between cedant and reinsurer, and also your views on the quantity and quality of the data you are using. (Think of this as a loading for "model risk" as opposed to the "insurance risk".) This loading might well vary from class to class and/or from submission to submission. This is entirely subjective, but loadings in the range of 10% up to 30% are not unheard of. (Your broker won't thank you for doing this but s/he doesn't pay your bonus...)
- It can be desirable to use a risk load that recognises the skewness of the loss distribution, and increases as you rise through a tower of cover (e.g. an increasing proportion of the standard deviation, which itself will also be increasing). Whilst this "double whammy" can generate very large percentage loadings to the (decreasing) expected loss cost, this is probably appropriate.
- In practice we believe many (most?) people in the London market are still using very simple risk loads.
 - For working layer casualty contracts this is often a simple grossing-up factor such as:
 - $\text{Required Premium} = 100 / 70 \times \text{Expected loss cost}$This is equivalent to setting a 70% target loss ratio (but don't forget to also consider investment income and internal expenses...). Typically the grossing up factor will reflect risk levels by varying by line of business and/or market (e.g. Professional Liability versus Auto Liability or US versus UK) according to the perceived differences in "riskiness". Again, these differentials will usually be quite subjective.
 - For property classes it is more common to see risk charges incorporated into the premiums as a percentage of the standard deviation of the modelled recoveries:
 - $\text{Required Premium} = \text{Expected loss cost} + 50\% \text{ of StDev}(\text{Expected loss cost})$Again, the factor (50% here) will usually vary by territory and peril (say 30% to 70% range in the less problematic zones?) and market conditions can blow these loadings sky high in certain conditions (e.g. Florida Wind is being priced at a somewhat higher level in the aftermath of the 2005 season).

The underwriters will often shift their loading parameters up and down in tune with the insurance cycle. For example, in the late 1990's a 90% loss ratio target for working casualty layers and a 30% factor for property reinsurance would not have been unusual, whereas currently we believe most underwriters are using significantly higher factors. This doesn't help the actuary looking for a stable benchmark against which to monitor market rate levels!

h – Rate increase analysis considerations

There are a number of ways of approaching the task of building a rate level index. The general methodology is usually to “rebase” the charged rates for the latest and previous exposure periods onto a common footing using the rating parameters (e.g. for 1 year term, \$1m limit, \$50k deductible, 1 exposure unit) and then compare these consistent rebased premiums to assess the change in rate charged per unit of exposure. Many of the issues and considerations are well documented in a 2001 GIRO workshop on “Premium Rating Indexes” chaired by Bill McConnell, the papers for which are available on the IoA website.

With reinsurance you will never get away from the fact that your index – assuming you even get one in the first place – will be (a) imperfect, (b) subjective or most likely (c) both of the above. As such, here are some things to consider when working with rate level indexes:

- The person making the subjective adjustments is usually the underwriter, who has a vested interest in demonstrating favourable rate actions. He or she will also have just written the risk and as such will be thinking good things about it. Whilst your underwriter undoubtedly approaches the task with due integrity, the end result is likely to contain some positive bias.
- If you are working with a rate level index prepared by someone else, it is vitally important to fully understand how it was compiled before you use it. For example, some people will wrap loss trend within their rate index to make a “profitability level” index, other people will keep trend out of the rate index and apply this separately.
- By virtue of the way they are compiled, rate indices are usually based solely upon analysis of renewal business. However, for some books of business there can be a high turnover rate and “new” business may comprise a significant part of the book.
 - Have the underwriters been offering slightly better terms to attract this new business onto the books? Or hopefully they actually charge *more* for the new business to reflect the greater risk inherent from the relative lack of knowledge of the new clients versus their renewal book... (Dream on!) The rate index will likely require some adjustment to reflect any differential in the general level of rates on new business versus renewal.
 - You can try to get a handle on this relativity by looking at the rates charged (rebased “per standard exposure unit”) for a block of new accounts compared to similar renewal accounts, but this can be difficult when every risk is “unique”...
- Also you shouldn't forget the accounts that were non-renewed or lost. How profitable were these compared to the renewed business? (We are venturing here into utopian areas where underwriters have access to and maintain detailed records on non-renewed risks etc, but even if you can't perform the numerical analysis this doesn't mean the consideration isn't an important one!)

- Consider this overly simplistic example:
 - \$100m of written premium in 2005 with an overall expected loss ratio of 70%
 - Homogeneous class of business with a nice steady 5% p.a. claims inflation
 - A premium rate study on business renewed into 2006 shows a 5% increase in rates
 - Written premium in 2006 is \$105m
 - This suggests 2006 loss ratio will also be 70%, right?

No, not necessarily... Suppose further that:

- \$25m of 2005 premiums are “good” with a 2005 expected loss ratio of 50%
- \$50m of 2005 premiums are “average” with a 2005 expected loss ratio of 70%
- \$25m of 2005 premiums are “poor” with a 2005 expected loss ratio of 90%
- Good and average risks were renewed at 5% rate increases: \$78.75m written premium
- Poor risks were non-renewed
- \$26.25m of “new good” risks were written, expected loss ratio 50%

For 2006 this means we are looking at...

- “old good” = \$26.25m premium and 50% expected loss ratio (rate and trend cancel)
- average = \$52.50m premium and 70% expected loss ratio (rate and trend cancel)
- poor = \$0m premium, all non-renewed
- new good = \$26.25m premium and 50% expected loss ratio
- Total 2006 = \$105m premium and **60%** expected loss ratio

Looking only at the headline “5% increase on renewal business” would have suggested a 70% loss ratio. Including the new business (only) would lead to an expected 65% loss ratio: \$78.75m of renewal premium at a “70%” loss ratio and \$26.25m of new premium at 50% loss ratio. It’s only when you consider all the components together that you get the correct picture.

Of course, the maths can (and frequently does!) work in the other direction too. Furthermore, working with real data is a mite less precise but you get the picture...

- Occasionally you might see analysis attempted to “retrospectively rate” new business in order to incorporate it within a rate review study, i.e. the underwriter (or actuary) reviews each new account and assigns a hypothetical expiring premium on the basis of what he/she thinks they might have written it for last year in the then prevailing conditions. BE AFRAID! Such analysis is both hugely subjective and dangerously biased!
- Another factor that can significantly skew rate level indices is what we would call “cycle spikes” i.e. chunks of opportunistic business written at appropriate times in the underwriting cycle. When markets are hard, the canny underwriter may write an extra-ordinary category of business at exceptional rates. When markets soften they will rapidly ditch this same block of contracts. Such behaviour can lead to exaggerated swings in the rate level index that are actually not reflective of what is happening in the “core” business. Applying this to your core book without adjustment can, for example, lead to poor a-priori estimates for a reserving Bornhuetter-Ferguson. We would recommend identifying and stripping opportunistic business out of the index, if you can...

- Stay alert to potential changes in the risks that are not captured by the rate level index. Any strengthening or weakening of terms and conditions will have an impact on the expected profitability beyond that explained by the rate index. What is the general market trend at the moment? Remember what we just said about bias...

i – ECO / XPL

“ECO” stands for “Extra-contractual Obligations” whereas “XPL” stands for “Excess of Policy Limits”. Both refer to (usually punitive) damages awarded by a court against an insurer above and beyond the coverage provided by their insurance policy, typically for bad faith, fraud or negligence when dealing with a claim. “XPL” relates to judgements over and above the policy limit on a valid claim whereas “ECO” relates to judgements that are beyond the normal scope of insurance coverage – for example, grossly negligent handling of an invalid claim might lead a judge to insist the insurer pay even though they would not have been liable had they handled the claim appropriately.

A couple of examples should help clarify the distinction.

- Suppose an athletics training club is successfully sued by one of its member athletes for a case of sexual abuse by one of the club staff. Under the terms of the club’s insurance policy sexual abuse is excluded, but insufficient supervision of staff is covered. If it could be argued that the sexual abuse arose because there was insufficient supervision of the staff then the claim would have two portions – the fact that the staff member was found guilty of sexually abusing the athlete (excluded under the policy) and the fact that there was insufficient supervision of staff in place (included under the policy). Theoretically the insurer should only pay the latter portion of the claim, but in practice the full amount could be paid to avoid legal disputes and to maintain goodwill with the insured. The extra amount is classified as “ECO”.
- A doctor is being sued for negligence. The doctor has a policy limit of \$1m and before the case goes to court, there is the possibility of an out-of-court settlement for \$750k. However, the insurer believes there is a good chance of winning and so the case goes to court. The final verdict is against the doctor and the award is for \$1.5m. At this stage, the doctor could appeal the decision but the insurer decides to pay the full \$1.5m to close off the case. The total payment made by the insurer has exceeded the \$1m policy limit held by the doctor, so this claim is “XPL”.

In reality there is a grey area between “ECO” and “XPL”. In the first example, the sexual abuse element of the claim might have increased the size of the claim beyond the original purchased limit, hence the claim would also be “XPL”. In the second example, the insurer was settling in excess of policy limits to avoid a claim of bad faith from the doctor because the claim had not been settled out of court (within the policy limit). Bad faith is not covered in the original policy, so it could be argued that this claim is partly “ECO”.

In practice, the vast majority of reinsurers provide “ECO” and “XPL” cover under the one banner “ECO/XPL” and price for it accordingly. This avoids a dispute in the future as to whether a claim is “ECO” or “XPL”.

ECO/XPL coverage can be provided in a specific reinsurance contract protecting only against such awards, or it can be included as an “add-on” within another reinsurance protection, as is frequently the case with clash covers.

There is some moral hazard here as the insurer's decisions in court might be influenced by having a reinsurance policy in his back pocket sweeping up all the "excessive" settlement. To mitigate this moral hazard, reinsurers typically require meaningful co-insurance participation by the direct insurer, so ECO/XPL claims are often only 90% covered with the remaining 10% reverting to the direct insurer.

ECO/XPL protections offer catastrophic type cover, i.e. they have (or should have!) a very small probability of a loss, but if there is a loss you know it's going to be a big one. As such they are generally priced by the market using a simple rate on line approach. For example \$10m of ECO/XPL cover might be sold for, say, 7.5% rate on line giving a premium of $0.075 \times \$10\text{m} = \750k with some kind of profit commission in the event of no recovery.

There's not a whole lot the actuary can do here, but we included our discussion of ECO/XPL for two simple reasons:

- To offer a simple explanation for anyone who hasn't come across it before;
- To note that this protection has a value and so should not be overlooked when pricing a reinsurance contract that includes ECO/XPL cover. At the very least it should be factored in to your rating process by charging a sensible rate on line based on how likely you think it might be triggered.

j – Cession treaties need pricing too!

Often when a reinsurance placement includes a cession treaty (usually the top layer of an excess of loss program) this layer will have little or no credible loss experience and will be priced purely on the basis of cession rates. The insurer will "pay as he goes", ceding a share of the premium on any risk with sufficient limit to expose the cession treaty. The share is determined according to a pre-agreed table of factors usually set out in the slip.

Cession rates are no more complicated than treaty allocation factors, or ILFs under another guise.

Often the broker will want to stop the "pesky meddling actuary" from getting involved as this could slow down his placement process. Said broker might suggest there is no point trying to analyse the cession layer because "it's never had a loss, there's no data for you to work with, it's money for nothing really, and you can't be selected against because you are getting a fair share of the premium on every risk as it is cession rated". You will not be surprised to hear that we disagree. The actuary can still add value:

- Are the cession factors based on ILFs that you feel are appropriate for the client and class of business – i.e. are you really getting (at least) a fair share?
- A fair share of what, exactly? Check the profitability of the gross premiums;
- A good loss record doesn't mean there is no meaningful exposure now, nor that you couldn't get tagged with a big loss! How exposed was the cession treaty in past years anyway? (The cession premium history should tell you if you don't have the historical limit profiles.)
- Keep an eye on those cession factors (or ILFs), especially in a softening market. Have they changed over time to your disadvantage? Most likely they have remained stable for years, but

this could well be to your disadvantage too as positive claims inflation may well mean the factors should trend upwards year on year.

- Make sure the calculation allows for any self insured retention or large policy deductibles. For example, suppose you have a £3m xs £2m cession layer and a £5m original policy limit. If the £5m policy has a £1m SIR it's really £5m xs £1m so rather than "3x2 part of a primary 5" you should be getting allocated "3x3 part of 5x1" which will be a much higher share. Some numbers to illustrate:

<u>Limit</u>	<u>ILF</u>
£1m	1.00
£2m	1.25
£3m	1.40
£4m	1.50
£5m	1.57
£6m	1.60

- Cession rate ignoring SIR = 3x2 part of 5 = $(1.57 - 1.25)/1.57 = 20.4\%$
- Cession rate allowing for SIR = 3x3 part of 5x1 = $(1.6 - 1.4)/(1.6 - 1.0) = 33.3\%$

Allowing correctly for the SIR results in a 63% higher reinsurance premium!

k – Don't ignore the words

We actuaries all love to get stuck into the numbers, but the words often contain qualitative information that is just as important, often more so. At a minimum we suggest you should always read the following parts of a reinsurance submission before finalising your analysis:

- **The slip.** This governs the contract and tells you exactly how things such as premium adjustments, currency conversions and so on will work. It also defines the important stuff like exactly who is being reinsured, for what classes of business, in what territories, with what exclusions, for what limits and retentions over which time period etc. Simply reading the slip should give you a firm understanding of the workings of the deal, and will often trigger questions you will want to ask your underwriter or broker.
- **The supporting narrative.** Yes, this can be tedious as they are often written by someone who has swallowed a superlatives dictionary and yes, they can be full of seemingly unimportant details that don't obviously affect your analysis. But they should give you some context and they often contain clues (or even blatant statements) to important dynamics that won't yet have come through the numbers. For example, a shift in emphasis of the book might be fundamentally important and might only really feature prominently in the narrative, especially if all the numerical exhibits are prepared looking back at the history.
- **The wording** (this might not be given to you automatically, but it should be available – especially with the dawn of "contract certainty"). These can be particularly tedious so you often won't want to read these cover to cover, but they include a lot of the key clauses you might want to look at specifically. For example, the chosen terrorism exclusion language.

With all of these wordy bits, it is often worth comparing new with old if you are renewing a risk and therefore have access to last year's documents. What has been changed, why has it been changed and how does it affect the terms of the contract?

Also, depending on how close you are to the final negotiations, it is worth checking the final slip against your final analysis to see if there were any last-minute changes such as the addition of a profit commission or other tweak that affects the expected profitability!

1 – Slip rates of exchange

Another reason to read the slip! When there are several currencies in play, it is common for a reinsurance slip to set out fixed rates of exchange to be used and these will not necessarily resemble prevailing foreign exchange rates. For example, in LMX casualty treaties you might see conversion fixed at, say, GBP1.00 = USD1.50 = CAD 2.40.

Fixed rates of exchange can also crop up sometimes in more detailed parts of a slip, for example within an Indexation Clause³.

All we would say here is that it doesn't really matter what exchange rates are used, *so long as* your analysis is performed on the same basis! Watch out for exhibits prepared on a different basis, such as consolidated sterling at the Lloyd's year-end rates. Depending on how significant the mix of currencies might be and how varied the profitability might be by currency, getting this wrong could have a significant impact on your conclusions.

³ See 5g of this paper for more on Indexation Clauses

4) DATA ISSUES

- a) No data – e.g. high excess layers
- b) When “As-If” becomes “If Only”...
- c) Tips for flushing out data issues
- d) Other practical miscellany

a – No Data, e.g. High excess layers

If you don't have any historical loss experience to work from you clearly need to use some kind of benchmark or industry-based exposure rating approach. The reasons should be obvious, but Pitfall 4 in the GIRO paper “1999 GISG Reinsurance Pricing Working Party” gives a numerical illustration of what will go wrong if you try to stick to experience rating.

Having said that, if you have a risk that has a long history with no losses, whilst this might not lend itself to an exposure rating exercise the good experience does tell you something! But do check the historical exposure levels to see how credible this “loss free period” really is.

So where does the aforementioned “benchmark” come from, then? There are no hard and fast rules, you simply have to do the best you can with whatever you can lay your hands on. This might be:

- A simple comparison with another similar risk (i.e. we charged X for a near identical protection to Client ABC so this client should probably pay something like X +/- a bit). You might call this “market rating” – but be careful you aren't writing a lot of similar risks all priced in line with something that was a bit dubious in the first place;
- Extrapolated curves based on the specific company experience that you do have (i.e. fit a severity curve to the losses in the lower layers and use this curve to impute rates in the higher layer). Don't be fooled into a false sense of accuracy - this is a very hit-and-miss process;
- Your own curves for the class of business based on your experience – this is probably the best alternative *IF* you have credible experience in the class. Note that territory can be just as important as class of business so don't automatically assume your Californian Lawyers ILFs will be appropriate for a book of Lawyers in DC – they probably won't be.
- Somebody else's curves based on “market wide” experience for the class. This can be better or worse than using your own curves. It should be based on a larger pool of data, but that pool might be somewhat less representative of your own particular risks if you are not writing a broad section of the entire market.

You also need to be particularly careful in creating a benchmark using data from a subscription market such as Lloyd's. Here the loss data from any one insurer will be *their share* of the loss based on the percentage line they signed on the original risks, which will vary from contract to contract – even within a program. This isn't rocket-science (when is it ever?) – you just need to be careful to remember what you are dealing with. *Intelligent execution...*

b – When “As If” becomes “If Only”...

A key step in the reinsurance pricing process is to “as-if” the data you are analysing. This means adjusting or stripping out losses to restate the contract experience “as if” the exposures of the past were more reflective of the current. Perhaps the most common examples are:

- An insurer has withdrawn from a class which previously had bad experience so they want you to strip this class out of the data on the basis that it will not be representative of current exposures. *“We couldn’t have those losses again, we no longer write that sort of business.”*
- They have significantly changed their limit profile – suppose they used to offer \$10m limits but now only write a maximum of \$5m. It is appropriate to strip out any loss dollars in excess of \$5m since *“we couldn’t get a loss that big in the coming year”*.

When performing any “as if” adjustment think carefully about what you are doing and why, and remember the inherent biases of those who are providing the information. We do not wish to suggest that anyone would knowingly tell you anything other than the truth, of course, but your typical cedant comes pre-programmed to work hard at looking after their own best interests. *Intelligent Execution* is fundamentally important to your view of the risk. Some things to consider include:

- The “As if” process is all too often a one-way street. You are usually being asked to consider removing all the nasty bits that would otherwise generate a high price. You are usually being fed information from a broker acting in good faith on what their client, the cedant, has told them, in search of the best deal. But what about any *new* exposures that may have been added to the contract in recent times and which as yet have generated no losses? What is the difference between stripping out some losses for removed exposures and *adding in* some losses in respect of these new exposures? Clearly it is easier to take out than to add in, and often you will be sold a line about how the new exposures are trivial (they were probably thrown in for free, after all...) but do you really believe this?
- Usually, losses are not the only thing that would have been avoided had they not done what they did back then. Losses usually have associated premiums so make sure that you try to adjust the premiums in a consistent manner. Don’t be surprised or put off when you are told this data is “not available”, just do the best you can. Doing nothing is probably not sensible.
- It is a glib statement, but there is some truth hidden in here: The worse the experience, the harder you will be pushed to as-if (some might say dream) away the losses. But whilst those specific losses might no longer be possible, other kinds of losses are. Those nasty old losses happened for reasons unforeseen at the time of writing – but there could be different (but equally unpleasant) unforeseen reasons now for future poor experience. What has *really* changed in the underwriting processes and controls? Are you fully convinced they really couldn’t make a similar mistake again? If you are, that’s great. But if you are sceptical, should you really be giving them full credit for the “as if” impact?

In extreme cases you might be asked to as-if away so many of the losses that you really don’t believe the result. For example, on a high layer excess of loss contract where you have few losses anyway it could be that as-if’ing away a class results in a loss free treaty. **“If only** we hadn’t written those contracts with the losses...”

One emerging issue stems from the increasing helpfulness of brokers (often in collaboration with us actuaries!) Most of the broking firms placing complex reinsurance contracts in London are these days adding more value to the process by pre-analysing much of the standard stuff that markets typically want to see and do. Of course this is not entirely new, but the prevalence and the sophistication of the pre-analysis is definitely on the increase and you will often be presented with data, exhibits and/or analysis that have already very helpfully been as-if'd in some way. Hopefully you will also be given the “raw” (unadjusted) information too so you can understand what they have done and why (and what the actual results were like!) but if you are not automatically given this, be sceptical and make sure you ask for it!

We are not saying that you shouldn't “as-if” your analysis. You should. We are also not saying that brokers present as-if analysis in order to mislead. Quite the opposite – they do it because they believe *not* doing so would be to mislead.

What we are saying is simply that the game is very definitely stacked against you so tread carefully, be rigorous, be internally-consistent and be ready to challenge!

c – Tips for flushing out data issues

You don't need us to tell you how important it is to work from good data, but often you will only be provided with a current snapshot of the exposure information together with a “slimline” version of the experience data (for example, only losses that breached 50% of the excess of loss retention). There will often be only a limited amount you can do to verify the data quality but that shouldn't prevent you from trying, particularly if you wrote the same contract last year...

- How do the actual losses compare to those you would have expected based upon the information provided last year? And how does this same picture look using data banded into layers such as \$100k, 150 x 100, 250 x 250, 500 x 500?
- Perform some simple comparisons of data reports from year to year. Premium volumes, splits by class or territory, policy limits and attachments, loss numbers and amounts, average premium per account, average signed lines, average exposure count, average policy term etc. Do they all seem consistent? Are there any apparent trends or patterns?
- Is there any facultative or other inuring reinsurance to consider? Has this already been allowed for in the data you have?
- Look at the loss frequencies and severities even if you plan on working with aggregate data as this can be very insightful. You may well have stable aggregate losses by underwriting year, but underlying this could be an increasing frequency trend concealed by a reducing average reported claim size. Is the increasing frequency a problem? The reducing claim size could well be a fallacy driven by differences in reporting delay by severity of loss, or by optimistic case reserving on less mature losses, or by something else...
- Have you been provided with a premium rate change analysis? If so, scrutinise this carefully as it is often the best window you will have onto the actual risk-by-risk underwriting that went on this year and therefore the expected change in profitability.

d – Other practical miscellany

- It is all too easy to underestimate the true loss frequencies for prospective exposure periods.
 - Beware the impact that loss severity trend has on loss frequencies for excess of loss layers! When analysing the experience for an excess of loss layer you would usually hope to have ground-up loss data for all losses bigger than, say, 50% of the layer attachment point. It is important to trend these ground-up losses *before* applying the reinsurance layer terms rather than the other way around, for reasons that should be obvious. If you do not have individual loss information and are instead working with only the aggregate losses to the layers, it is important to note that you are not only missing some information on loss severities due to trend, but you are also missing some loss frequency detail.
 - Furthermore, with long tail classes IBNR is always a factor and this is especially true with excess of loss reinsurance due to (a) the additional time it takes for a loss to penetrate the attachment, and (b) the additional links in the communication chain.
 - Understand your reporting pattern! Does the reinsurance respond on a “losses occurring during” or a “claims made during” basis? What about the underlying policies? What is the full chain of events from accident to reinsurance recovery, and how does this shape your loss reporting tail? Have there been any significant changes in this chain in recent years that might distort the loss development, such as Tort Reform (i.e. a change in insurance law) or a shift between claims made and occurrence policy forms? If you don’t know, ask your underwriter...
- When adjusting claims frequencies on-level to current exposures, take care including historical years with small exposures. For example, suppose you currently have an annual exposure base of 500 units (say \$500m of on-level premium or 500 class-1 physicians or whatever else it might be) and 10 years ago the client suffered 3 losses from an exposure base of 50 units. Simply pro-rating this up to 30 losses against “today’s exposures” could be very misleading. That 30 could very easily have been 20 or 40. It may be better just to exclude years where there is little credibility to the estimate, or use some kind of credibility weighted average across the years?
- Changes in limit / attachment profile can have an impact on a cedant’s loss experience. At the simplest level, a gradually increasing trend in limits offered from year to year will mean that the historical loss experience may be suppressed relative to current exposure. In practice this can be difficult to handle, not least because you frequently don’t know the past limit profiles. However, one interesting approach to this problem can be found in the 2005 paper “An improved method for experience rating excess of loss treaties using exposure rating techniques” by Mata and Verheyen.
- Think about the “big picture”. What is happening with your premium volume, average limit and attachment, premium rates, policy count etc and what does this imply? For example, if you are told that rates are flat and premium volume is stable, but the average policy limit has increased significantly, has the policy count gone down (bigger risks, more severity, less frequency?) or stayed the same or even increased (something doesn’t stack up... have the underlying rates really decreased or has there been a shift in the book?)

- Be careful to use a stable exposure measure, in particular note that written premium volume is usually far from stable! If you only have written premiums to work with, these should at least be adjusted on-level with current premium rate levels in order to cancel the distortions due to rate fluctuations. (See also section 3h of this note on the dangers of using such rate indices blindly!)
- Watch out for under-reserving (or, less commonly, over-reserving) distorting your data. This is especially common on ALAE, where a simple analysis of case averages (paid versus reserved) will often reveal apparently optimistic ALAE reserves.
- Trend your historical losses, but be careful not to over-trend! Compound interest is a powerful beast and it's all too easy to generate unreasonably high trended losses. Do your trended results look reasonable? Will your underwriter believe them, or are you risking your own credibility by presenting something that leads him to conclude you've taken leave of your senses?
- Recognise that trend is rarely well quantified, rarely a simple compound rate and that "unlimited" claims inflation (i.e. on the full loss value before applying any insurance or reinsurance limits and deductibles) can be a significantly different number to the claims inflation relevant to your layer. The following overly simple example makes this issue obvious:
 - \$2m average ground up loss, apply 6% inflation for 1 year = 2,120,000
 - For a \$4m xs \$1m reinsurance policy, trend would be 12% ($1120 / 1000$)
 - For a \$1m xs \$1m reinsurance policy, trend would be 0% ($1000 / 1000$)

There are a number of old (but still valid) papers readily available for a more detailed study of trend. The following ones give a pretty good overview of the key issues:

- Adjusting size of loss distributions for trend (Rosenberg, 1981)
- The effect of trend on excess of loss coverages (Keatinge, 1989)
- How to choose a trend factor (Krakowski, 1994)

5) PRICING “BELLS AND WHISTLES”

- a) General comments
- b) Limits on Cover: AADs, Corridors, Mandatory Co RI, Aggregate Caps
- c) Retrospective Premium Adjustments: Swing Rates, RPs, APs, PCs, NCBs
- d) Franchises and Reverse Franchises
- e) Top & Drops
- f) Multi-year vs annual
- g) Indexation clauses

a – General comments

This section considers a number of commonly encountered reinsurance contract features that in some way modify either the recoveries or the total premium ceded and describes how you might handle them when assessing technical rates and prices. Please remember, we are talking “traditional” reinsurance and have not considered the way in which similar or other features might be incorporated into an alternative risk transfer (ART) reinsurance product. Furthermore, potential accounting and reserving implications (of which there are several⁴) are beyond the scope of our paper.

Some of the existing papers do already cover some of this in brief terms, but much of this is based on using an aggregate loss distribution which these days might be unusual. (You might well have a frequency / severity model or you might simply be using deterministic methods but how often do you go half-way?) We felt it warranted covering again here because it is such an important part of practical reinsurance pricing. You almost never see a reinsurance contract that doesn’t have some kind of bell or whistle attached to it, and occasionally you get a full orchestra!

We illustrate some of our points here using pricing measures and techniques. We are not suggesting that either the methods or the loadings used here are appropriate to any given line of business or situation; we are just using them to illustrate what sort of thing you might do.

Much of this section (indeed the whole paper) is common sense and all falls under the “intelligent execution” banner, i.e. think carefully about how the reinsurance treaty will work and make sure what you are doing reflects that. In particular, it is vital that you fully understand the interaction of all the bells and whistles and apply them in your modelling in the correct order. For example, applying a large retention and then an aggregate cap will produce a very different result to doing it the other way around. The slip should always make it clear, or you can phone a friend(ly broker)...

The basic approach to all of these bells and whistles is fundamentally the same. We make some generic comments here on the use of frequency / severity models, experience rating (i.e. burning cost type calculations) and exposure rating (using ILF curves etc). We will then look at each bell or whistle in turn and make any specific comments as need be, which are primarily describing how they usually work in practice!

⁴ Actuaries interested in reinsurance accounting might find it useful to review the US accounting rules (FASB 113 for overall guidelines and also the supplementary note EITF 93-06 for multi-year contracts) whilst those with a reserving focus will find that many of the pricing techniques described can be adapted for assessing reinsurance recoverables at early development periods on long tailed classes.

- *Frequency / severity simulation*

This is arguably the most reliable approach to pricing bells and whistles – and also, incidentally, for demonstrating risk transfer to your auditors – simply because it allows you the flexibility to model and quantify virtually anything. The hardest part (other than setting your frequency and severity assumptions) is simply making sure you understand how the contract works and then building that logic into your analysis.

A second benefit of this approach (over exposure rating and experience rating) is that you can see and incorporate the impact on volatility as well as the expected values. The impacts on volatility can be very significant so there can be direct implications for the size of risk load you might want to use.

If you are performing a net present value (NPV) assessment, note also that many of these twists (for example the aggregate deductible and all the premium adjustments) will affect the cash flow patterns as well as the nominal amounts.

- *Experience Rating (Burning cost)*

As with all burning cost analyses, prior to applying historic losses to the layer, these must be trended for claims inflation, developed for IBNR / IBNER and adjusted on-level for changes in exposure between the loss year.

The on-levelling stage can be critical with some of these features, particularly aggregate retentions and the like. A small error in on-levelling can get magnified into a big difference in the indicated rates.

In practice the burning cost method can be very hard to use effectively to assess many of the bells and whistles described below.

- *Exposure rating*

Exposure curves essentially provide technical rates for reinsurance with unlimited, free reinstatements and no aggregate deductible or other bell / whistle. You will need to make some kind of adjustment based on your judgement. Again, this can be difficult and you will usually resort to a frequency / severity approach.

One approach to using exposure (or experience) rating with bells and whistles is to price the “bare” contract as normal, but then to look at the relationship of either the expected loss cost or your indicated premium rate based on simulation, with and without the particular feature. You might then adjust the experience / exposure rate for either the absolute difference in these two, or based on the ratio of one to the other.

When trying to quantify such contract features it is important to remember that they are usually there for good commercial reasons. The market price adjustment for the addition of such a feature may therefore differ significantly from the theoretical impact on the technical rate. Possible reasons include:

- Additional comfort from cedant sharing the loss experience = less need for safety margins?
- Some markets (i.e. reinsurers) will not write unlimited covers, so by restricting the coverage you attract more underwriters and increase competitive pressure on the rates;
- In a particular reinsurer, capital allocation may favour certain features;
- Market cycle effects – e.g. NCB commonly offered in soft cat market – perhaps gives the underwriter the illusion he is not reducing rates! In a hard market limiting features may be forced on cedants with no rate reduction etc.

This should not stop you trying to quantify the bells and whistles, but it does give you something else to think about!

b – Limits on Cover: Aggregate Deductibles, Corridors, Co-Reinsurance, Caps

Many (almost all, these days?) excess of loss reinsurance programmes do not offer the cedant unlimited, free reinstatements. You will usually find one or more features restricting the overall aggregate limit available. Examples include aggregate deductibles (often called “annual aggregate deductibles” or “AAD”), loss corridors, mandatory coinsurance and aggregate recovery caps.

It is important to make allowance for these restrictions when making any assessment of reinsurance, and to assess their impact on the effectiveness of the reinsurance in meeting the cedant’s risk mitigation requirements.

Aggregate deductibles (as applied to excess of loss reinsurance)

An “aggregate deductible” (sometimes called an “AAD” – shorthand for ‘annual aggregate deductible’ – or often an “inner aggregate”) is best described by example: A reinsurance contract may cover losses up to \$2.5m any one loss in excess of \$2.5m any one loss, in excess of \$5m in the aggregate (or in broker speak, “2½ x 2½ x 5”). Here the individual recoveries (\$2½m xs \$2½m) are calculated as normal, but the first \$5m-worth of these recoveries is retained by the cedant within the aggregate deductible.

AADs are typically found on so-called “working layers” where you would otherwise anticipate a high level of loss activity. The aggregate retention removes an element of dollar-swapping between cedant and reinsurer and reduces the expected recoveries and hence reinsurance premium.

Commonly (but by no means exclusively) the AAD is set at a level that still allows a relatively high likelihood of loss activity in the reinsurance contract. “Back-up” layers are sometimes also purchased to provide additional sideways coverage behind a reinsurance layer with a limited number of reinstatements. For example, if the main reinsurance layer was \$4m xs \$1m with an aggregate loss cap of \$20m the cedant might also buy a “backup” layer of \$4m xs \$1m xs \$20m. Clearly this is just a large aggregate deductible again so warrants a similar approach, but remember the context – backup layers usually have little credible experience and require particular care when pricing.

Loss corridors (as applied to quota share or excess of loss reinsurance)

There is currently some debate in the US as to whether a loss corridor effectively bifurcates a reinsurance contract into effectively two separate contracts, only one of which may involve risk transfer and hence qualify as reinsurance. We have not attempted to cover risk transfer or accounting issues within this paper, so here we simply discuss the way in which this feature can be modelled.

Loss corridors are more typically seen on quota share than excess of loss contracts. The loss corridor represents an aggregate deductible for claims that would otherwise be recoverable under the treaty. However, unlike an AAD, a loss corridor is placed between two tranches of cover within the same layer so the reinsurer pays some losses before the corridor bites.

For example, suppose we have an 80% quota-share on a £12.5m book of business. The quota share might be placed with benefit of a loss corridor from, say, 85% to 100%. This would mean that the cedant would recover 80% of his claims until his loss ratio hit 85%, and a further 80% of claims for that part of the loss ratio exceeding 100%. Between 85% and 100% he makes no reinsurance recoveries. Assuming (as is usually the case) the reinsurer received the full 80% share of the premium, and ignoring any brokerage or commission, if the cedant had a gross loss ratio of 107% the application of the loss corridor in this example would result in a cedant net loss ratio of 167% whilst the reinsurer enjoyed a loss ratio of 92%:

<i>£000's</i>	<i>Premiums</i>	<i>Losses</i>	<i>Ratio</i>
Gross	12,500	13,375	107%
80% (before impact of corridor)	10,000	10,700	107%
Corridor attaches (80% terms)	10,000	8,500	85%
Corridor exits (80% terms)	10,000	10,000	100%
Reinsurer = 80% net of corridor	10,000	9,200	92%
Net to Cedant	2,500	4,175	167%

The equivalent for an excess of loss contract would be to have a reinsurance contract of, say, \$2.5m xs \$2.5m with a loss corridor of \$5m xs \$7.5m. Here the reinsurer would protect the first three limits of recovery, and the sixth and subsequent limits of recovery.

The allowance for loss corridors in technical rates for loss corridors is essentially undertaken in a similar way to that described for aggregate deductibles. Where the loss corridor applies to a quota share, one could use either a stochastic loss ratio simulation method or an adjusted burning cost method to calculate the gross loss ratio and recoveries before and after application of loss corridor.

Mandatory co-reinsurance

Reinsurers may sometimes include a requirement that the cedant retains a proportion of the reinsurance layer themselves by way of co-reinsurance (co-ri) which is “net and unreinsured”. The cedant will not be permitted to reinsure out this co-ri by any other mechanism.

In modelling terms, co-ri can simply be treated as a part-placement. It should, in theory, have no impact on the rate charged. For example, suppose the rate for a layer of \$900,000 xs \$100,000 is 23% of Subject Net Premium Income (SNPI). If the contract carries a requirement for a 10% co-ri, then the ceded premium is simply calculated as

23% x SNPI x 90%

In practice, however, commercial considerations can come into play. A cedant's willingness to retain an element of co-ri can be seen as a very positive sign by reinsurance markets, and consequently may result in a slightly better rate or a more enthusiastic take-up by reinsurance markets where previously there had been a placement shortfall.

Aggregate recovery caps

Aggregate recovery caps might be expressed in monetary terms (e.g. \$7,500,000), as a percentage of ceded premium (e.g. 400% of ceded premium), or as a number of reinstatements for the layer (e.g. \$2.5m xs \$2.5m with 2 reinstatements). All three are simply different methods of expressing the same aggregate cap on recoveries.

Where reinsurance contracts offer different sections of coverage under a single contract it is not unusual to see a separate aggregate limit on each section plus an overall aggregate limit across the entire contract. For example a contract with three sections might have the following aggregate limits:

Section A	\$10m
Section B	\$10m
Section C	\$5m
All sections combined	\$15m

If you are modelling each Section separately it should be easy to allow for this, but if not you might simply cut a few corners such as applying the \$15m cap to the overall but ignoring the effect of the sub-limits if you believe that the error will not be significant. (Of course, this will reduce your modelled expected recoveries a little. This could be a prudent or imprudent approximation, depending on whether you are buying or selling.)

c – Retrospective Premium Adjustments: Swing Rates, RPs, APs, PCs, NCBs

There are a number of retrospective, loss sensitive premium adjustments regularly included within reinsurance contracts. These include swing rates, reinstatement premiums (RP), loss dependent adjustment premiums (AP), profit commissions (PC) and no claims bonuses (NCB). Full allowance for these features must be included when estimating the technical rate for a reinsurance contract.

Swing rates

There are two main types of swing rate formula, which we will call a “pure” basis and a “minimum-plus” basis for ease of explanation:

- **Pure:** A reinsurance contract may express the premium rate as a provisional rate with minimum and maximum rates and a “swing” of some percentage of recoveries, typically in the range 50 - 125%. For example a contract may have a provisional rate of 10%, a minimum of 4% and a maximum of 20%, and a swing of 120% of recoveries. In this example the premium that will actually, ultimately be charged is 120% of ultimate recoveries, but subject to this amount being no less than the minimum rate of 4% and no more than the maximum rate of 20%. Initially the provisional premium is paid but then this is adjusted at stated intervals until all losses are settled. (Read The Slip!!)

- **Minimum-plus:** Here you see the same terminology used (provisional, minimum and maximum rates, and a “loss load”) but they work somewhat differently. (Read The Slip again!!) Here for example we might have a 4% minimum, a 10% provisional, an 18% maximum and a 110% loss load. Again the premium paid initially is based on the provisional rate, but as losses develop the premium is adjusted to:

$$\begin{aligned}\text{Swing Rate} &= \text{lesser of } \{ \text{maximum rate, [minimum rate + (recoveries x loss load)] } \} \\ &= \text{lesser of } \{ 18\% \text{ of GP, } 4\% \text{ of GP} + 110\% \text{ aggregate recoveries } \}\end{aligned}$$

where GP = gross subject premium income.

When swing rates are applied to long-tailed business, the slip will usually lay out all the practical aspects of the premium adjustments, which usually start some time after the end of the reinsurance year, and sometimes incorporate a formulaic IBNR / IBNER factor. The IBNR formula does not usually impact the ultimate premium paid, at least in nominal terms, as the annual adjustments continue until all losses have been settled.

Swing rated contracts may be useful where the cedant and the reinsurance market have strong opposing views of the likely recoveries, for example after an account has been re-underwritten following heavy losses, or for a volatile class of business, or for a new account with no track record. Compared to a fixed premium alternative, the swing can allow the cedant to pay a lower premium if their loss experience turns out to be as benign as they claim, and the reinsurers get a larger premium if the loss experience is adverse.

Swing rated contracts have a lower volatility in reinsurer’s profit / loss than the equivalent contract written at a fixed rate because they effectively transfer less risk. When the premium is still swinging (i.e. until the losses are high enough to max it out) the two parties are dollar-swapping and the reinsurer is taking a margin to cover the risk that he carries in the event the swing reaches the maximum and the losses continue to rise.

You may be asked to propose an appropriate swing rate structure that meets the reinsurer’s risk, capital and profitability requirements. This is not an exact science but possible approaches include:

- An iterative simulation method: Estimate a “first guess” swing rate structure and then refine this through a number of simulations, each time adjusting the swing structure to move the distribution of reinsurer profit / loss closer to that required.
- An alternative is to dump each iteration of the simulated aggregate losses into a spreadsheet. You can then parameterise the swing factors in your spreadsheet and play with them (maybe using a solver-type function) to determine appropriate structures. Although this may be less cumbersome than the iterative simulation method, in practice the calculation method may take longer due to calculation speed and simulation output file size.

Reasonable places to start for the “first guess” could be based on reviewing the summary percentiles of aggregate recoveries, or by looking at the flat exposure rate for the treaty.

Reinstatement premiums

A reinstatement premium may be charged for reinstating the cover following a reinsurance loss payment. For example, “one reinstatement” means that following exhaustion of the reinsurance limit you can reinstate one additional limit. “Two reinstatements” means you can do this twice, etc. So a \$5m limit with 2 reinstatements means that you actually have three lots of \$5m – so \$15m – aggregate cover, but you may have to pay extra premium – the RP – after the first and second limits are exhausted.

Reinstatement premiums are expressed as a percentage of the basic ceded premium; a contract may have different percentages for different reinstatements, for example 2 free, 2 at 50% and 1 at 100%. In this example, the cedant would not pay any reinstatement premium for the first two limits of coverage used, they would pay 50% of the basic premium for each of the next two limits used (on a pro-rata basis if the full limit is not used), and would pay 100% for the next limit. The final (sixth) limit would not carry any reinstatement premium.

Reinstatement premiums reduce both the net expected loss cost (i.e. net of reinstatement premium) and the volatility in this net loss cost. For reinsurance layers where high levels of activity might be expected, the inclusion of reinstatement premiums can be highly significant to the technical rate.

(Note that brokerage is not usually payable on reinstatement premiums. Again, read the slip!)

For calculating a reinstatement structure, it may be most useful to concentrate initially on what the cedant / reinsurer aims to achieve, for example reduced costs in an average year, or reduced volatility in an adverse year, and then use the simulation model to refine this basic structure.

Also, if using a simulation model to compare RP terms – think carefully about whether you are happy using the typical Poisson frequency model given it’s relative lack of skew. We spoke about this earlier (§3c) but it is pertinent here as this skew affects the value of the RPs.

Loss dependent adjustment premiums (AP’s)

Loss dependent AP’s are based on the recoveries under the contract. They can be expressed either as a percentage of basic premium or, more commonly, as a percentage of the recovery. There could be a flat rate AP across all recoveries, but it would be more usual to have different AP levels for different tranches of aggregate recoveries. An example is shown below:

AP as % recovery	Recovery band
0%	\$0 – 5m
15%	\$5m – 10m
25%	\$10m – 20m
5%	\$20m – 30m

The approach described above for pricing contracts with paid reinstatement premiums also applies to loss dependent AP’s. One minor difference is that brokerage may be payable in AP’s – this varies by contract. (What should you read? That’s right.)

Profit commissions (PC's)

PCs are essentially a means of returning a portion of the ceded premium to the cedant in the event of favourable loss experience. The calculation depends entirely on how these are specified in the reinsurance contract. At their most simple they may take the form:

$$\text{Experience account} = (A\% \times \text{ceded premium}) - \text{recoveries}$$

$$\text{PC} = \text{Max} (0, B\% \times \text{experience account})$$

where A and B are less than 100. $(100\% - A\%)$ reflects that part of the premium which goes to pay reinsurers' expenses, and B% reflects a sharing in the profit between reinsured and reinsurer.

On longer tail business lines it would be more usual to calculate an experience account that includes investment income and makes specific allowance for the timing of cashflows. Furthermore, profit commission clauses will often include a deficit carry forward, which may be either limited in time or unlimited. In this case, if the experience account is negative, this negative value is carried forward into the next year's PC calculation and offsets the first part of any profit in that year.

Unless the reinsurance contract also carries a loss share provision, the cedant has no obligation to pay any experience account deficit back; however it will be offset against profits in future years before the PC is calculated.

Two final points of note on PC's:

- In cases where the experience account includes a credit for investment income, this rate might be somewhat less than the rate used to discount cashflows for profitability testing, for example the experience fund credit rate could be set at LIBOR – 50 bps.
- The way in which the PC is triggered / paid varies widely. In some cases it is payable only if the contract is commuted; in others a provisional payment is made initially and then adjusted annually thereafter, with the calculation based on a predetermined methodology of estimating IBNR; with some contracts any overpaid PC can not be refunded but only be offset against PC's on future years – which you may not have written...

The variation in PC calculations means it is essential to carefully read the PC clause in the slip and wording for every contract modelled with this feature. Furthermore, your modelling approach for PCs may differ depending on whether or not there is a deficit carry forward clause:

- No deficit carry forward: It is a simple matter to create the experience account and calculate the PC based on the simulated claims when the premium is already set and the contract simply needs to be tested for profitability. When calculating the premium and / or PC structure, the calculation becomes recursive and must be solved either by use of a solver-type function or by successive simulations each with a refined estimate of premium and PC.
- With deficit carry forward: If the deficit from past years can be well estimated, e.g. for property catastrophe reinsurance, then this estimate can simply be included in the calculation described above. For longer tail classes, it is generally necessary to include deficit carry forward calculations from older years based on deterministic estimates, but to include a stochastic element for more recent past years. This requirement to model on a stochastic multi-year basis on both past and future exposure periods means you must give careful consideration to the degree of correlation within the class between years.

No claims bonuses (NCB's)

No claims bonuses are essentially a predetermined partial refund of premium on commutation in the event that there have been no recoveries under the contract. Typically there is a restricted period after the end of coverage, for example 60 or 90 days, during which this option can be exercised. This feature makes NCB's more suitable for short tail classes, and they are most commonly seen on property catastrophe excess of loss reinsurance.

NCB's are either paid in full or not at all: Recovering just \$1 under the contract will void the NCB. This can easily be factored in to your analysis – simply calculate the minimum value of a claim that would be worth making and “throw away” simulated losses below this amount. However, in reality it could well be that a kind of “gentleman’s agreement” might be followed whereby an underwriter would treat the NCB like a PC (i.e. give partial credit) if the claims were “low”. This could be influenced by the underwriting cycle, such “generosity” being more likely in soft market conditions than hard.

d – Franchises / Warranties and Reverse Franchises

Franchises can be considered to be an “on/off” switch for the reinsurance – if the condition of the franchise is not met for a particular loss or event, then no recovery can be made in that instance. Franchises can be applied either to the size of the market loss, or the size of the loss to the cedant.

Probably the most familiar market loss franchises are Industry Loss Warranties (ILW's) for catastrophic event and Original Loss Warranties (OLW's) for aviation losses. In the case of an ILW, the total loss reported by the specified body (such as PCS for US perils) must exceed a given value to trigger the reinsurance. Once triggered, the reinsurance is applied in the normal manner.

One may also see franchises applied to the size of loss to the cedant. For example, a cedant might purchase a layer of \$500,000 xs \$500,000 with a franchise of \$5m. In this case, only losses exceeding \$5m would be applied to the layer. The reinsurer is protected from excessive loss activity in the layer caused by a high frequency of small losses. Meanwhile the cedant has additional coverage on larger losses; this might be beneficial if higher layers have retrospective loss sensitive premium features (such as paid reinstatements), or if these layers have not been completely placed (co-ri or shortfall).

Similarly, you may see reinsurance layers with either a reverse franchise, or both franchise and reverse franchise. A reverse franchise requires that the loss or event is *less* than a pre-determined amount. The reverse franchise may be applied to reported market loss for catastrophic events, or to the loss size to the individual cedant. For example, a market level reverse franchise could specify that the reinsurance would apply only for events for which the market loss reported by PCS was less than \$10bn. A cedant level reserve franchise might specify that only losses less than \$5m to the cedant qualify for the reinsurance layer.

Where a reverse franchise is combined with a franchise (sometimes termed a franchise corridor), a reinsurance contract will only trigger if the loss or event is between two pre-determined values. For example, the reinsurance might protect for events where the Sigma reported loss exceeds \$2.5bn but is less than \$10bn.

Frequency / Severity Comments

There are a number of ways in which one can apply franchises within frequency / severity modelling. Modelling this feature is easier if the franchise applies to the cedant loss rather than the market level loss:

- *Franchise on cedant loss value:* If there is a single, simple franchise, you have only to total the recoveries for the applicable simulated losses. Alternatively, adding a single cell to each loss with a 1 / 0 indicator (1 for reinsurance switched on, 0 for reinsurance switched off) allows you to address franchise corridors or franchises that vary by type of loss, peril or territory.
- *Franchise on industry / original loss value:* One aspect that must be borne in mind is that if the franchise is based on industry or original loss size, for each event you must simulate both the size of the original event and the loss to the cedant. For cat modelled property this is relatively straight-forward provided you have access to both the portfolio event loss table (ELT) and the industry loss curves (ILC's) – although even here you must determine the level of correlation between the size of market loss for a specific event and that for the portfolio if you wish to include the secondary uncertainty element available from RMS ELTs.

For other classes and perils, it might not be possible to construct a market loss curve; for each historic loss the actuary might know whether or not the franchise would be triggered, but not the actual size of the market loss. An option here is to generate a uniform (0,1) value for each simulated loss. You must then create either a formula or a lookup table that identifies the probability that a loss of that size (or in that loss size band) will trigger the market loss franchise. If the value simulated from the uniform distribution for that loss is lower than this probability, the franchise is assumed to be triggered.

If you can generate a distribution of market losses, the difficulty remaining is in calculating the cedant loss that is appropriate for each simulated market loss. One approach you might use is to estimate the “miss factor”, representing those market losses where the cedant has no or minimal exposure and so cannot experience a loss large enough to be recoverable under the reinsurance. For the remaining losses, you could then generate a market share distribution.

- For example, the cedant loss as a percentage of market loss might be expressed as a lognormal distribution with mean 5% and standard deviation 4%. For each simulated market loss you now simulate two additional figures – one being a discrete distribution with values 1 / 0, to indicate whether there is any loss to the cedant, the second being the cedant loss as a percentage of market share. These are then applied to the market loss value to generate the cedant loss. The final step is to cap off the cedant loss at the maximum line size if this is a “per loss” reinsurance rather than “per event”.

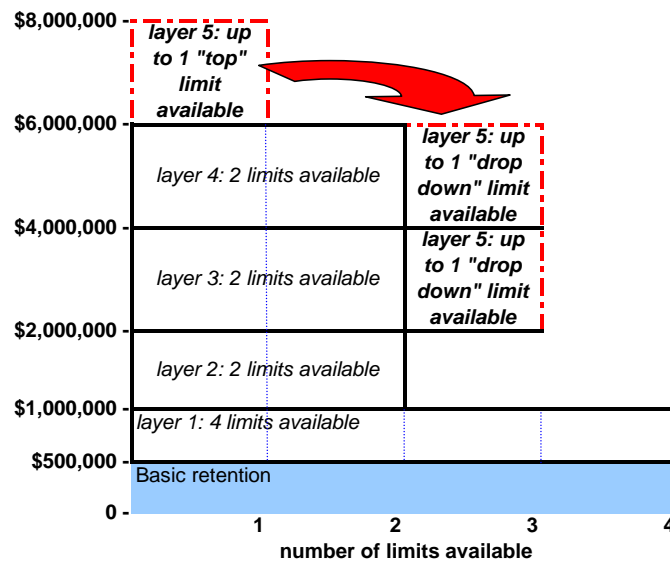
Experience Rating Comments

If the franchise is based on a market loss value, the actual market loss associated with each historic loss must also be adjusted on-level. Review of the adjusted market loss will then indicate whether the adjusted historic loss should be applied to reinsurance.

e – Top & Drops

A “Top & Drop” reinsurance contract offers protection that can either be applied as a high (typically the highest) reinsurance layer stacking vertically on top of other layers, or that can be applied as a back-up layer to one or more of the lower layers of protection.

The diagram below shows an example of a top & drop contract. The top & drop is layer 5 in this example. It provides \$2,000,000 of cover in all. This can be used to provide cover in the layer \$2m xs \$8m, or as a back-up to layers 3 and 4 (acting as \$2m xs \$2m xs \$4m in the aggregate and \$2m xs \$4m xs \$2m in the aggregate respectively).



In practice, Top & Drop contracts may offer reinstatements (paid or free). Additionally, only part of the overall aggregate limit may be available to drop-down. Top & Drop contracts often incorporate a “no higher, no deeper” warrant from the cedant that they will not purchase a layers higher than the “top” section or sitting sideways behind the “drop” section.

Frequency Severity Comments

Probably the easiest way to model a top and drop is to model this initially as two or more sections, and then to limit the aggregate recoveries across all sections. The “top” section of the layer is modelled as a normal top layer. The “drop” section is modelled separately as a back-up contract for each layer it sits behind. Taking the contract above as an example, the sections of the contract would be modelled as:

- “top” - \$2m xs \$6m
- “drop-down (a)” - \$2m xs \$4m xs \$4m in the aggregate
- “drop-down (b)” - \$2m xs \$2m xs \$4m in the aggregate

The total recoveries are then calculated as the lesser of the sum of the three sections or the total aggregate limit available (in this example \$2m).

Some care must be taken with any NPV assessment to ensure the cashflows included are those that allow the cedant to make recoveries at the earliest possible opportunity, even if the internal allocation of recoveries to specific losses might differ.

Exposure Rating Comments

If you are confident the cedant is unlikely to have more than one limit loss at the top of the programme, then the exposure rate for a layer at this level with unlimited free reinstatements should differ very little from a “one-shot” contract. Consequently, you might accept the exposure rate as a reasonable rate for the “top” element of the cover.

The “drop-down” element is more problematic, as it would be very difficult to adjust a “straight-in” exposure rate to reflect the difference in price for a back-up layer with a far more remote likelihood of activation. You might look at the percentage increase in rates on the frequency / severity or burning cost model when comparing a (capped) top layer with the rate for a top & drop with the same aggregate limit. This increase could then be applied to the exposure rate. Alternatively you might consider the probability of the “drop” being triggered, assume it would be a limit loss in this case and charge a simple rate on line.

Note that the exposure rate might be more appropriate for the “top” element of the contract, as the cedant’s historic experience of very large loss might not be reflective of their actual risk.

f – Multi-year vs. annual

Multi-year contracts are a lot less commonplace now than they have been, driven largely by regulation in the Alternative Risk arena making it increasingly hard to structure an effective multi-year contract that reinsurers are willing to sell and cedants are keen to buy. However, they do still exist and they throw up some interesting considerations.

The first thing to note is that there are a number of potential accounting issues⁵ that can significantly impact the operation and therefore effectiveness of a multi-year reinsurance contract, and these issues vary by territory. Accounting considerations are beyond the scope of this paper but we strongly recommend that anyone investigating a multi-year contract should seek advice from their finance department and, if appropriate, their auditors.

Focussing on some of the pricing issues in modelling a multi-year contract, we suggest you consider the following:

- The aggregate limit over the full term may well be less than the sum of the annual limits; this aspect can be modelled in a way similar to that outlined for top and drop contracts above.
- Correlations between consecutive years – these can be strong and should therefore be modelled;

⁵ For example, you may need to accrue profit / loss shares and a proportion of future years’ premium if loss experience this year could impact the potential recoveries or premiums in future years. For an overview of the types of features that can cause accounting issues on multi-year reinsurance in the US in particular we suggest reading EITF 93-06, which you should find here: <http://www.fasb.org/pdf/abs93-6.pdf>.

- Underwriting or Accident year coverage basis? Which is more appropriate depends on the risks being protected. For example, for property catastrophe a multi-year contract is more effective on an accident year (LOD) basis, as a single event could impact two or three separate underwriting years. For classes driven more by underwriting conditions, underwriting year coverage might be more effective. (Note that in either case you should consider whether a new case law, regulation or business practice might impact multiple years);
- The socio-economic environment and the way in which this might change over the term of the multi-year contract. You might consider a time-series based economic model that is then correlated to claims and recoveries, for example;
- The relative strength of the downgrade clause needs to be considered: For multi-year contracts the downgrade clause may differ significantly from those in more traditional one-year contracts. For example, if the reinsurer is downgraded below A-minus it may be required to collateralise the unpaid element of the remaining limit. when assessing multi-year contracts. Furthermore, given the multi-year nature, the clause is made more important because the credit risks are increased by the longer tail and the larger aggregate numbers;

The above list is not exhaustive, but it gives a taste of what you are likely to find in multi-year reinsurance programmes.

With all of these it is important to bear in mind that the market adjustment for the addition of such a feature may differ significantly from the impact on the technical rate. There are a numerous possible commercial reasons for such differences as previously discussed, such as:

- Comfort taken from cedant sharing in the loss experience;
- Restricting coverage can open up new markets and increase competitive pressures;
- In a particular reinsurer, capital allocation may favour certain features;
- Market cycle effects.

g – Indexation Clauses

Many liability reinsurance contracts, especially motor reinsurance contracts, are subject to indexation clauses in respect of bodily injury claims. These clauses serve to limit the inflation risk transferred to the reinsurers by the long term nature of the recoveries. This is achieved by increasing the retention (and the limit) by the amount of inflation over the period until all recoveries are settled.

The basic form of an indexation clause would see the retention and limit indexed with reference to a named inflation index, e.g. UK wage inflation, to the time of each payment of each claim. It is an option, albeit less common, that the retention and limit are indexed to the final settlement date of the claim.

There are 3 main types of clause used in common practice:

- **Full Indexation:** This clause transfers all of the inflation risk to the insured. The retention and limit are indexed from day one to the date of actual payment by the insurer.
- **Severe Inflation Cause (SIC):** The SIC provides a mechanism to share the inflation risk between the insurer and reinsurer. There is no adjustment to the retention or limit until the increase in the reference index exceeds a defined percentage, e.g. 30%. Once past the defined percentage point the retention and limit are only adjusted for increases above the 30%.
- **Franchise Inflation Cause (FIC):** The FIC provides another mechanism to share the inflation risk between the insurer and reinsurer. There is no adjustment to the retention or limit until the increase in the reference index exceeds a defined percentage, e.g. 30%. However, once past the defined percentage point the retention and limit are adjusted for the full increases from day one of the contract.

Modelling the indexation clause

To correctly handle the indexation clause, the amount and time of each payment of each claim should be modelled. However, the most common approach is to apply an average increase to the layer across all claims. This method indexes the retention and limit with reference to the mean term of payment and expected future inflation for the given index.

$$\text{Adjusted Retention} = \text{Retention} \times (1 + \text{inf})^{(\text{mean term})}$$

$$\text{Adjusted Limit} = \text{Limit} \times (1 + \text{inf})^{(\text{mean term})}$$

The mean term of payment can easily be calculated from the payment pattern. The adjusted retention and limit can then be used in the burning cost, exposure rating or frequency/severity models as required. The limitation of this method is that it does not allow for the distribution of settlement periods and may underestimate the full impact of the clause.

Given the geared nature of reinsurance recoveries it is essential that an allowance is made for the indexation clause in the loss cost calculations. The following example demonstrates the impact the indexation clause has on XL pricing:

Example

- No of claims = 1 per year
- Severity Dist = Lognormal with $\mu = 12.97$, $\sigma = 1.2$, Threshold 1m
- Future Inflation = 5%
- Mean term until settlement = 6 years
- Layer = 2m xs 2m

Indexation Method	Indexed Layer	Expected Loss Cost	Saving vs. No Indexation
No Indexation	2m xs 2m	472,021	n/a
Full Indexation	2.68m xs 2.68m	380,454	20%
SIC 20%	2.28m xs 2.28m	431,226	9%
FIC 20%	2.68m xs 2.68m	380,454	20%

Note that the FIC effectively switches from the “No Indexation” method to the “Full Indexation” method once the indexation threshold is exceeded. As such its value should always lie between these two other methods, the relative position determined by the height of the indexation franchise.

6) UNDERWRITERS' RULES OF THUMB – HOW TO TEST THEM

- a) 1x1 as rate on 1x0, 1x2 as rate on 1x1, etc
- b) Discounts for aggregate caps and limited / paid reinstatements
- c) Minimum Rates on Line

Throughout the reinsurance industry you will encounter underwriters using apparently un-scientific rules of thumbs to assist in the pricing of various reinsurance layers or contract features. It can be tempting to dismiss these as without technical foundation and therefore somehow “wrong” (especially if the underwriter cannot provide any objective justification) but don’t be so hasty – many such benchmarks have evolved over generations of underwriting experience and wisdom, and some are based on a sound footing, it’s just that the technical justification has been lost in the hazy mists of time and too many underwriter / broker lunches...

On the other hand, sometimes these rules may have descended from something like relativities based upon, say, a certain pareto curve. The cumulative effect of loss trend over the intervening decades could well have served to invalidate such relativities.

Rules of thumb can usually be tested for reasonableness with a little actuarial thinking and a simple spreadsheet, as we will now illustrate. Our list is obviously far from exhaustive, we merely seek to demonstrate how such tests might be approached.

a – 1x1 as rate on 1x0, 1x2 as rate on 1x1, etc

All the numbers presented here are purely illustrative and should not be used for real pricing!

- Example Marine rule of thumb:
 - Price up a base layer of £10m
 - Each subsequent £10m layer priced at 80% of the immediately underlying layer price
 - So if £10m limit is worth £100k, then £10 xs 10m is £80k, £10 xs 20m is £64k, etc
- Example PL rules of thumb:
 - Price a £1m primary limit
 - £1m xs £1m = 30% of primary £1m
 - £3m xs £2m = 80% of £1m xs £1m
 - £5m xs £5m = 100% of £3m xs £2m price
- Example EL rules of thumb:
 - Price a £10m primary limit
 - £10m xs £10m = 15% of primary premium.

These kinds of rules have largely been replaced now by the explicit use of ILFs (Increased Limits Factors) but you might still encounter the old style rules occasionally. Such rules can be quickly compared against an ILF table for the relevant class of business, or by reviewing an appropriate loss severity curve if you have one.

For example, it is easy to deduce that the rules of thumb given above are consistent with the following ILFs:

<u>Marine ILF</u>		<u>PL ILF</u>		<u>EL ILF</u>	
£10m	1.00	£1m	1.00	£10m	1.00
£20m	1.80	£2m	1.30	£20m	1.15
		£5m	1.54		
		£10m	1.78		

All you need do is convert the rule into a set of ILFs and then consider whether the ILFs look sensible.

b – Discounts for aggregate cap or limited / paid reinstatements

It is very common during the placement of a reinsurance treaty to encounter eleventh-hour negotiations between broker and underwriter, horse-trading the price against the depth of coverage afforded (i.e. the aggregate limit) and any loss sensitive premium adjustments such as an RP (reinstatement premium) or AP (additional premium). Get this part of the negotiation wrong and it is all too easy to undo the hard work you have done in generating a robust base price by giving away too much of your expected underwriting margin.

- Example 1: 10% discount given against the unlimited exposure price in return for a 200% aggregate loss ratio cap;
- Example 2: Contract is priced at a 20% rate on line assuming unlimited free reinstatements. A 20% discount is given (charged rate 16%) in return for a limit of one paid reinstatement at 100% of the up-front premium (i.e. “1 @ 100”);
- Example 3: Contract is priced up assuming unlimited free reinstatements. 20% discount given in return for a limit of two paid reinstatements, each being charged at 50% of the up-front premium (i.e. “2 @ 50”).

Usually the underwriters will be quite happy using their experience and judgement to assess this trade once they are happy with their base price. However, this doesn’t mean to say the actuary cannot add value at this stage in the negotiations by offering a real-time, objective comparison of whatever alternatives the underwriter / broker / reinsured might dream up between them. A simple spreadsheet is all you usually need.

Taking our simple examples to illustrate the point...

Example 1 = 10% discount for 200% loss ratio cap

Here you can choose an appropriate aggregate loss distribution to fit around your expected recoveries (you may already have one, depending on the analysis you have performed). Make sure that you are happy with the spread and shape of this distribution before you use it, for example by considering the 5th %ile and 95th %ile against the best and worst loss ratios experienced by the account in the last 10 years (adjusted on level for rate and trend etc). You can then use this to quantify the discount.

Suppose your expected losses were 80% of the premium (i.e. an 80% loss ratio) before the imposition of a loss ratio cap, and you had chosen a LogNormal distribution with parameters $\mu = -0.4$, $\sigma = 0.6$. Using this you can quickly find the expected value of the limited distribution using @Risk or other such software:

$$E[\text{LogNormal}(-0.4, 0.6)] = 0.8055$$

$$E[\min\{\text{LogNormal}(-0.4, 0.6), 200\%\}] = 0.7820$$

So, the cap reduces your expected loss ratio from 80.6% to 78.2%, a saving of only 3%. As such a premium discount of 10% looks too generous – based on your assumptions... Giving the 10% discount would see your modelled loss ratio expectation rise from 80% to $0.8 \times 0.97 / 0.90 = 86\%$.

Take care when performing such back of the envelope calculations – your conclusion are only as good as your models and assumptions.

Example 2 – X% RoL Unlimited => X% discount for 1 @ 100.

For example, suppose the contract is rated at 10% rate on line (i.e. premium divided by policy limit = 0.1, e.g. £1m premium and £10m limit) on a simple flat rated and unlimited basis, i.e. a fixed premium and no constraint on the number of losses you could recover under the treaty.

This rule of thumb would suggest a 9% rate on line (i.e. $10\% \times (1 - 10\%)$) for an otherwise identical contract with coverage limited to 1 paid reinstatement at 100% RP, i.e. you pay 9% up front, then pay a further 9% (being 100% of the original premium) in the event of a full recovery, and only two full limit recoveries are allowed.

In fact, this example is built upon the simple assumptions that all recoveries are full limit losses (no partial recoveries) and that the loss frequency follows a poisson distribution. Whilst these two assumptions might often be challenged (especially the first one) they are usually reasonable approximations to the truth for low rate-on-line deals...

A little algebra is all you need to show that the “9%, 1 @ 100” is reasonably equitable to the “10% flat” alternative. First off, given the low rate on line and poisson assumption, we can say that in broad terms the number of losses will be either:

- 0 with c.90% probability
- 1 with c.10% probability (10% rate on line)
- 2+ with c.0% probability

This means that the aggregate cap on recoveries has no significant impact, i.e. both treaties pay the same expected losses. It also means that both treaties pay almost the same expected premiums as the RP basis pays either:

- 9% with c.90% probability
- 18% with c.10% probability

Expected premium = 9.9% vs flat rated 10%.

However, there are a couple of flaws here. Okay, so there are a lot of flaws here... but there are two that are particularly noteworthy:

- The rule of thumb becomes less reliable as the rate on line (loss frequency) increases;
- The above calculation ignores the fact that a 10% rate on line probably means a 5% or less loss frequency as such contracts will be priced to a high expected profit margin.

Looking at the second of these for a moment we can see that this shouldn't be a show-stopper. Suppose we have a 10% rate on line still and are therefore going to offer a 10% discount, but the loss frequency is actually only 5% ("pure" contract priced to an expected loss ratio of 50%). Switching to the 1@100 basis again doesn't materially affect the recoveries, but expected premiums are reduced:

9% with c.95% probability
18% with c.5% probability

Expected premium = 9.45% vs flat rated 10%.
Expected loss ratio = 53%, i.e. $50\% \times 10 / 9.45$

Here the expected loss ratio has crept up from 50% to 53%, but this should still be okay when considering firstly the uncertainties in that original 50% anyway and secondly the reduced aggregate exposure through the limited reinstatements.

Remember, changing the shape of the contract has implications for the required supporting capital and as such a small increase in loss ratio when you do something that reduces your worst case scenarios may well be entirely reasonable and still consistent with your profit goals.

Example 3 is a straightforward extension of Example 2 and as such can be tested in a similar way. Alternatively, you can build a simple @Risk spreadsheet to validate the premium structure. (Using a simple simulation model can also be a way to allow for the possibility of partial recoveries...)

c – Minimum Rates on Line

These are commonly imposed on very high / remote reinsurance layers, i.e. where the likelihood of making any recovery is very slim. Here the expected loss cost might be near zero but the reinsurers will not price the risk on this basis, they want a decent sum of money. Why do cedants still buy these layers when the prices are out of kilter with the expected losses? Because they pretty much have to. Just because your model doesn't have any losses at that level doesn't mean there is no exposure, and the typical cedant simply cannot afford to have a net loss at that level. Two examples could be:

- The unlimited liability component of a UK motor portfolio
- "PML Failure" on a property cat program (i.e. the risk that your modelling has underestimated your true exposures and you blow through the top of your reinsurance program – Katrina gives us a recent example of how this might happen...)

The minimum rates used move all the time based on market conditions, usually without any apparent statistical basis and often justified on the basis of "our minimum rate to cover the cost of capital used

supporting this risk”. However, in reality there is often no such calculation behind the scenes, it is just the underwriter imposing their selected minimum pricing.

However, there could still be some merit in validating the minimum rate on line against your internal cost of capital. This will depend upon how you allocate capital at your company and also your target returns, but on a very simple level you might say something like:

- The reinsurance contract is hardly exposed at all, my assumed loss frequency might be 1%
- Capital might be allocated on the basis of, say, 99% VaR. Ignoring any diversification benefit, this would suggest a capital requirement in the order of the policy limit (effectively 100% VaR). There will presumably be some diversification benefit so maybe we assume capital of 75% of a limit??
- Your “cost of capital” might be interpreted as your target underwriting profit as a percentage of capital, which let’s suppose for illustration purposes is 12%.
- Minimum rate on line for this contract should therefore be 10%, because the 10% premium would just cover the 1% mean loss cost and the 12% required profit on capital (remember, you just notionally allocated capital equal to 75% of policy limit: $12\% \text{ on } 75\% = 9\% \text{ on } 100\%$).

Yes this is crude, yes this is dirty – but at least it gives you some kind of rationale for the minimum rate on line. Ideally you would refine this approach to something more technically robust based on your own capital models.

Don’t be surprised if when you perform such checks you occasionally find that a layer does not provide the required returns even using the minimum rate on line, suggesting that the minimum should be higher. This is quite common. However, before you start jumping up and down and waving flags at your underwriter who wants to write this contract at these prices, it is also important to remember the context. Usually such high layers will form but one part of a reinsurance program that is being placed simultaneously so there may well be some degree of cross-subsidy. Brokers are usually very canny at ensuring markets only write the “juicy” layers if they also take a share on the less juicy ones within a program. Does your intended participation across the program make sense?

On the other hand, it is also not unusual to encounter layers where the minimum RoL being demanded by the lead underwriter would appear to be over the top. But remember, at this level your position is highly leveraged and very uncertain so a large risk load is appropriate, as is a healthy profit load.

7) COMMON PITFALLS

- a) Ignoring drift in underlying limit / deductible profile
- b) Using ILFs without considering currency, trend, LAE, etc
- c) Pricing Retro treaties on ventilated programs without “looking through”
- d) Trend issues
- e) Extrapolating past years of a fast growing account
- f) Ignoring “Section B” (e.g. clash or other “incidental” exposure)
- g) Changes in underlying coverage e.g. Claims Made vs Occurrence

a - Ignoring drift in underlying limit / deductible profile

One of the main potential pitfalls of using experience rating (or frequency/severity fitting) techniques is that changes in the underlying exposure are not taken into account unless specific adjustments are made for them. As discussed previously, changes such as historical pricing, overall exposure and inflation will normally be taken into account.

There can be, however, some more subtle changes in exposure that take longer to manifest and be recognised in the observed claims data. Good examples of this are changes in the underlying limit / deductible profile.

Clearly a decrease in original policy deductibles means that the likelihood of a claim affecting the reinsurer is increased. Also an increase in an original policy limit meant that higher reinsurance layers can now possibly be affected.

This pitfall can be a particular issue when rates in the original market are falling. There are two main reasons why. Firstly, the original cedants might choose to take advantage of the falling market by buying cover lower down for the same premium. Secondly, the insurer/reinsured, in an effort to maintain premium volume, might write higher limits and/or lower attachment points than previously considered.

This can give rise to a “double-whammy” since, as well as the exposure increasing, the business may be being written at lower or inadequate premium rates. It is in such original market conditions that there is most pressure on the reinsurer to reduce rates as well.

Although the converse situation (the failure to recognise an upward drift in deductibles or a downward drift in limits) might not be so worrying from a conservative pricing point of view, it does mean that business might be declined unnecessarily.

b - Using ILFs without considering currency, trend, LAE, etc

If you looked by year at total annual losses and total annual losses above a threshold X, say \$1m, you would not be surprised to see that the ratio of losses above the threshold to the total losses increases due to inflationary effects.

However, what is less often appreciated is that the situation with ILFs is no different. As they are also based on monetary amounts the ILFs also need adjusting over time to reflect inflationary trends and failure to do so will result in under pricing excess layers. The issue is apparent when you look at the

way ILFs are usually constructed. The details are beyond the scope of our paper⁶, but essentially it involves inflating historical loss data up to today's values and calculating the factors from this rebased data. Clearly if the original loss data had been inflated by one more year (to 2007 rather than 2006 for example) the calculated factors would have been different.

Another potential pitfall with ILFs is applying those from one country to another country, perhaps due to lack of available ILFs or data in this second country. Other than the obvious issues about how applicable the ILF is likely to be in another country (e.g. they may have a very different legal climate with direct consequences for loss severity curves), you also need to consider how the ILF should be adjusted for the exchange rate. For example, using £1 = \$1.50 for illustration, if the ILF for a \$2m limit against a \$1m limit is 1.4, this should probably be taken as the factor for a £1.33m limit vs a £0.67m limit, and not as the factor for a £2m limit against a £1m limit.

Another related ILF pitfall is not understanding the LAE basis, and not ensuring that this is consistent with the layers you are pricing. For a worked example of this pitfall see section 3e of this paper.

c - Pricing Retro treaties on ventilated programs without “looking through”

There are various scenarios where the risk being reinsured is not completely transparent. Let's consider a couple of examples:

Example 1

A reinsurer writes a £3m xs £2m risk excess of loss layer on two PI insurers A & B

- Insurer A writes a mix of primary and excess business with a gapping strategy (sometimes called a “ventilated” program). On a particular risk they might write a 25% share of both a £10m primary and a 12.5% share of a £20m xs £20m excess layer, making their total exposure to the risk £5m.
- Insurer B writes only primary business. On the same risk they might write a 50% share of the £10m primary layer, making their total exposure to the risk £5m.

In one sense each insurer's exposure to this risk might look the same to the reinsurer “£5m total exposure attaching excess £0.” However, let's calculate what size an original insured loss needs to be to give the reinsurer a loss of £1m on this risk. A £1m loss to reinsurer will in each case be a £3m loss to the insurer since the reinsurance layer is £3m xs £2m.

A £3m loss to Insurer A would require a £24m ground-up loss to the original insured [25% of 10 + 12.5% of 4 gives 3, i.e. the original loss needs to go £4m into the £20m xs £20m layer]. A £3m loss to Insurer B however will only be £6m to the original insured [50% of 6 is 3].

Clearly to price this reinsurance correctly, the reinsurer needs to understand all the insured exposures at the individual insured risk level. This is “looking-through”.

⁶ See “On the theory of increased limits and excess of loss pricing” by Miccolis, 1977

Example 2

Similar considerations apply on certain types of property business. Consider a commercial property worth £100m and insured by two insurers, each taking a 50% quota share of the full £100m.

- Insurer A buys facultative cover for £30m xs £20m on their £50m exposure, and also has a treaty reinsurance programme for £15m xs £5m.
- Insurer B also buys facultative cover for £40m xs £10m on their £50m exposure, but has a treaty reinsurance programme for 75% of £10m (net of fac).

Suppose a reinsurer covers both insurers, writing a £5m line on each of the facultative policies and also taking a £5m share of each of the treaties. The reinsurer therefore has four separate £5m exposures but evidently they are all on the same underlying risk and can therefore respond to the same loss events. Furthermore, the way in which these £5m recoveries will develop relative to the underlying exposure will be different.

If the reinsurer purchases retrocession, all these sources of exposure to the same risk can accumulate and give a potential loss of £20m to the retrocessionaire. Equally, each £5m slug has a different “value” to the retrocessionaire.

These two examples show that when pricing reinsurance or retrocession, the reinsurer / retrocessionaire ideally can benefit greatly from “looking through” to the underlying risk details to better understand the exposure they are taking on. This means obtaining and interrogating the individual risk bordereau and not pricing off the back of a limit / attachment profile...

There is often however, insufficient information available to enable this, particularly for retrocession where there are 2 intervening parties to look through.

d - Trend issues

Trend issues are particularly important when dealing with reinsurance, and especially with excess of loss layers. The trend assumption is very often one of the most important (and most subjective) assumptions in the pricing process.

Things to consider include:

- It is likely that an excess of loss layer will have a higher trend rate than the underlying rate due to the geared nature of the retention. Furthermore, the effect of policy limits can make the converse true of primary layers and even some low-level excess of loss layers.⁷
- One trend almost always allowed for is for claims inflation. The appropriate rates to use here will differ by class of business because of the differing drivers of the inflation.
- It is less commonly recognised that claims inflation can vary by claim severity, with larger losses often exhibiting higher inflation.

⁷ See Keatinge's 1989 paper “The effect of trend on excess of loss coverages” for more detailed illustrations.

- Trends are also not necessarily constant over time. Certain ‘shocks’ or one-off events can cause a jump in claims severity. For example, consider the impact of Ogden discount rate changes on UK motor. Trend factors calculated without adjusting for these will include an implicit allowance for future ‘shocks’ as well as the general trend. This may be considered an appropriate thing to do, but care must be taken that when adjusting for known future ‘shocks’, that the trend is not being partly double counted.
- A change in mix of business (or limit/deductible profile) can exhibit itself in an apparent trend in the loss history. In this case it is also important not to double count the change, once explicitly and once implicitly through the trend analysis.
- Remember also that the frequency of losses (per unit exposure) may also exhibit a trend as well.
- Another difficulty with trends is the deciding of whether a historical ‘trend’ is really a trend or just good/bad luck. For example, if you look at the number of fatal accidents from US airlines the data from the last 5 years is much better than the preceding 5 years. Is this real safety improvement or is it just “good luck”?
- It is usually the case with reinsurance that you do not have enough contract specific data from which to impute a trend factor.
- Have a look at “How to choose a trend factor” by Krakowski, 1994.

e – Extrapolating past years of a fast growing account

When an account is growing fast it is important when pricing to understand exactly why. Questions might include -

- Is the insurer undercutting the rest of the market and putting on business at uneconomic rates?
- Are they offering improved terms and conditions and, if so, are these properly rated for in the original rates and incorporated into their rate index?
- Are they writing business that other insurers have stopped writing?
- Have they changed their underwriting guidelines and are now writing types of risk not previously accepted?

Even if the growth is coming from similar business to the existing, you still need to be aware of the fact that the low data volume from the past business could have low credibility and that applying an exposure based growth factor to the past loss frequency could materially under or over estimate the true loss frequency.

f – Ignoring section B (e.g. Clash / Incidental exposure)

Another pitfall is ignoring (or not paying enough attention to) section B type coverages (i.e. additional reinsurance exposures that might be deemed “incidental” or “non-core” and thus feature in a second section of the slip.

The exact type of coverage given here tends to vary more than the main section. This can be for a variety of reasons, but one of the potential larger impacts is the position in the insurance cycle. In softening markets, section B coverage may well be increased by insurers trying to remain competitive without reducing the top-line premium too much. It’s a neat way to soften the terms away from the headline numbers that underwriters have to report to their management. Conversely, in hard markets, coverages like these are sometimes the first to be attacked.

The “add-on” nature of this coverage may lead to it being neglected in the pricing process. This is particularly the case when there is little or no historical exposure or loss data about these exposures. Even when data is available, the way the coverage may have varied in the past may make the appropriate use of this historical data difficult or impossible.

These “Section B” exposures will tend to be of lower frequency and more volatile than “Section A” so are harder to model even if good data was available. The cyclical nature of this coverage and its use as a negotiating point means that “Section B” coverage will tend to be specific to each company. This prevents the application of any general loading.

Read the slip fully and recognise that *all* risk transfer has a value and a capital implication!

g – Changes in underlying coverage e.g. Claims Made/Losses Occurring During

The basis of coverage (CM or LOD) is fundamental to pricing. Development patterns will be very different, and profitability (ULRs) is very likely to be impacted too. Also, the exposure measure used for a CM basis needs an additional allowance for past volumes as the current policy year will be exposed to past problems.

A change from CM to LOD coverage or (more commonly) vice versa creates an overlap or a gap in coverage that needs to be addressed. It is important in such cases to understand any time limits on a previous coverage basis and how they will interact going forward.

Where there is a gap in coverage, or where the insured stops writing business reinsured on a CM basis, separate run-off cover may be bought or in some cases included as a reporting extension to the CM policy to allow for such eventualities.

Remember, it is also possible that it was some element of the past experience was part of the reason for a change in coverage basis. For example, a run of poor experience in a liability class written on an LOD basis will often lead underwriters to demand a switch to CM in order to gain more rapid control over the situation (witness US Medical Malpractice in the 1980’s). In such cases it is likely that there were other important changes to coverage, exposures and/or underwriting criteria that will also need allowing for. Your underwriter should be able to tell you more about this...

8) ACCUMULATIONS

- a) Pricing Clash
- b) Correlations for the Real World
- c) Systemic Losses

a - Pricing Clash

Due to the limited amount of data available, it is often quite difficult to accurately price a Clash cover. Actuarial judgment *and* market knowledge are necessary so you should discuss the exposures with your underwriter. Sometimes the source of a potential clash will be obvious (e.g. medical malpractice, where you often have several doctors and a hospital all named in the same lawsuit) and in other cases it might be less obvious (e.g. a Lloyd's syndicate writing a diverse book of international property and casualty business).

In “working” clash layers – i.e. where there is regular loss activity – experience and exposure rating techniques may work well, in particular the use of a simple stochastic model to simulate the number of clash events and the number and severity of each component within the clash. Such models are pretty intuitive to set up and explain, and can cater for all the nuances of a given clash layer such as aggregate retentions and limits. At the other end of the spectrum (“unexpected clash”) you will likely be in the realms of a simple rate on line.

But before you start playing with numbers you should take time to ensure you fully understand what you are working with. How does the contract work? Where are the exposures? Exactly what data you have been provided with?

- How does the clash protection interact with other reinsurance layers? Often you will see clash provided as a “section B” on a working layer excess of loss contract. For example, a treaty might provide \$750k xs \$250k each and every loss, and also \$1,250k xs \$250k on a clash basis. This would usually work such that the 750 xs 250 recoveries are made first on each loss irrespective of whether it forms part of a clash event, and then the 250k retentions on each component would be added together for each clash event and applied to the 1250 xs 250 layer. Read the slip – it should make it clear!
- How does the insurer monitor its accumulations? Sometimes this simply is not possible, for example Marine Hull and Cargo classes obviously clash but it isn't feasible to monitor the contents and exact whereabouts of every container throughout the year. If the accumulations are not properly monitored, do not rely exclusively on the available information as it probably underestimates the risk. Build in a margin for the unreported accumulations.
- If you have been given a risk profile (summary data such as a limits profile, territorial split etc) has this prepared on a “per Risk” basis (i.e. summing policies across any one insured) or “per Policy”? Have all the accumulations been aggregated in the risk profile (e.g. Building *and* Contents) or have you just been given the largest individual component? This will impact your view on the likely severity of any given clash event – i.e. are you being shown the components or the aggregates?
- If (as is likely) your experience and exposure rating uses mostly non-clash claims, you need to exercise caution using the same factors and curves for clash. For example, if you were working from large loss data for the “each loss” layers, the fitted curves are likely to overstate the average severity of a clash component as clash events often sweep up small items too.

- When frequency / severity modelling a clash event it is important to make sure your frequency and severity choices are fitted to consistent data. For example, you might have a class (such as medical malpractice) where your cedant's clash experience shows that a typical event has 5 or 6 components, but that usually 3 of these components are trivial in value. If you use a severity curve fitted to non-trivial components only you need to use a frequency measure also based on non-trivial components or you will significantly overprice.
- The price should be tested against reasonable scenarios with given return periods – this is the way many underwriters will reasonability-test prices. The rate on line offers a quick rule of thumb for doing so. For instance, a 5% RoL suggests a return period of more than 20 years. (At a 100% expected loss ratio it would imply 20 years, at a 50% expected loss ratio it would imply 40 years, etc.)
- In the case of a Lloyd's syndicate it is often worth referring to their Lloyd's RDS (Realistic Disaster Scenarios) numbers to get a feel for their view of the "worst case" scenarios – in particular the Liability Risks RDS and the Terrorism RDS will each involve clash events.

b - Correlations for the Real World


Correlation is omnipresent in reinsurance pricing. It always has an impact... but how significant is it? It could be anywhere from negligible to dramatic, depending on the context. It is quite commonplace to see correlations ignored in pricing on the basis that "they are too hard". But another way of thinking about this head-in-sand philosophy is that you are simply making an assumption of zero correlation coefficients. Are you happy making this assumption?

It is a valid point that correlations can be hard to set up and especially to parameterise as you rarely have the necessary data. But that shouldn't stop you having a go! We would suggest keeping it simple and therefore explainable...

- Before modelling correlation, consider whether it is likely to have a significant impact on your results. If you really think not, don't waste too much time on this – concentrate your efforts on other parts of the modelling;
- Correlation between loss frequency and severity (i.e. whenever you get a lot of losses, they also tend to be big ones) will influence the value of aggregate deductibles or reinstatement provisions;
- Think what the key correlations might be and try to model these explicitly. For example, if you are pricing a whole-account treaty that covers six lines of business, two of which are correlated to, say, stock market performance – in your simulation model you can hook these together.
 - One simple way to start is to generate a random variable "x" modelling the likelihood of, say, a significant stock market crash – this could be as simple as a Bernoulli trial with a sensible success probability – if and only if this triggers in your simulation you manipulate your correlated risk sources in some way (e.g. sample from a different, more severe distribution?)

- Due to the lack of data, it can be difficult to model correlation rigorously. Even the covariance of 2 variables is often difficult to guesstimate. It may be that a simple model of linear correlations makes more practical sense than a more complex method (copulas etc). The time saved with the simplified modelling can be used to test the model;
- Try correlating a minimal number of variables with crude correlation coefficients, such as by restricting your choice of correlation to one of:
 - 0% No correlation
 - 30% Weak correlation
 - 60% Reasonable correlation
 - 90% Strong correlation

When doing this, remember that 30% correlation for example really is weak and will not have a dramatic impact. The following picture taken from Wikipedia⁸ illustrates this rather nicely:



- It is also worth noting that it will be a lot quicker and easier to produce your desired skew by using a small number of well correlated factors than it will by using a large, more complex web of factors;

⁸ <http://en.wikipedia.org>

- “The Positive Definite Issue”: Any correlation matrix you define needs to be “positive definite” in order to make sense. In layman terms, if A is correlated to B and B is correlated to C this puts a constraint on the possible correlations between A and C so you need to make sure this is consistent. If you don’t like the thought of the maths, don’t worry – there are algorithms around that can help, and if you are using a DFA model the chances are such an algorithm will be built in. @Risk also has a function for testing for “PD-ness”.
- Any more on this subject would be too technical for this paper. Have a look on the internet if you want to scratch a technical itch, e.g. try <http://www2.gsu.edu/~mkteer/npdmatri.html>.

c - Systemic risk

The exact definition of a “systemic risk” or a “systemic loss” is open to a bit of interpretation as there is no hard and fast dividing line between one of these and, for example, a large clash loss. How big does a large clash have to be before it becomes a systemic problem? However, let’s leave that debate to the lawyers for now. For our purposes a “systemic event” is one that generates a market-wide unexpected accumulation of claims, potentially over a multi-year period, but usually triggered by a calendar year effect such as a stock-market collapse or the findings of a regulatory investigation. Recent examples include the so-called laddering losses in the US and endowment mis-selling in the UK.

Systemic events pose a particular threat to professional liability accounts, and can often be closely correlated with specific market practices or prevailing economic conditions. Recession or stock market collapse for example can drive behaviour that leads to the unearthing of systemic problems.


The other practical issue we have with systemic events is that once they are known, insurance and reinsurance terms and conditions very quickly evolve to exclude future reoccurrences and regulation also seeks to prevent them happening again. You will invariably be asked to “as if” these events out of the experience (see section 4b of this paper) and you will naturally be inclined to do so, but stopping there would under-price the treaty. What about the next systemic event? You can’t know what it is, when it might happen or how big it might be but you can be sure the exposure is out there and (depending on the class of business) it needs to at least be actively considered in developing your pricing models.

- The simplest (and arguably best) approach is to strip the systemic events out of your data to develop your loss cost for “regular” losses but then to calculate and add back a loading for the systemic events. This is easy to incorporate into a stochastic model:
 - Simulate a random indicator that is “on” with a certain frequency, i.e. your chosen return period for a systemic event. You will need to pick this based on a review of how many such events have affected your chosen class / territory in the past, combined with your view of current loss propensity relative to the long run average. For example, if there have been 3 events in the last 20 years (15% frequency) but you think the current environment is ripe for another event, you should choose a higher probability;


- When this indicator is on, you frequency / severity generate a block of systemic losses and add them to the general losses. The severity pattern of these losses might well be steeper than your regular losses, with policy limits possibly coming heavily into play. Note that modelling the frequency this way is much better than simply increasing the overall loss frequency a little as you model the clustering effect that can be so damaging to ignore;
- When choosing your “component count” parameter (i.e. the number of policies being hit by an event once it is triggered) the on-level premium is not an ideal exposure measure to use. It can help to look at the policy count now versus historically, and in particular to think about the number of risks that might conceivably be hit by one systemic event.

For US securities class action law suits (which include many of the past systemic issues) the Stanford Law School has a great website⁹ containing a wealth of freely available and useful information:

- *“The Securities Class Action Clearinghouse provides detailed information relating to the prosecution, defence and settlement of federal class action securities fraud litigation. The Clearinghouse maintains an Index of Filings of 2431 issuers that have been named in federal class action securities fraud lawsuits since passage of the Private Securities Litigation Reform Act of 1995. The Clearinghouse also contains copies of more than 16,000 complaints, briefs, filings, and other litigation-related materials filed in these cases.”*



Securities Class Action Litigation by Type of Lawsuit



⁹ <http://securities.stanford.edu/>

9) CREDIT RISK ISSUES

- a) Pricing approaches
- b) Considerations
- c) Modelling credit risks using CDS prices

Although reinsurance transfers insurance risk away from the insurer, it also transfers credit risk onto the insurer i.e. the risk that the reinsurer will not be around to pay claims in the long term. This credit risk can be very large – the reinsurance recoverable is often the single largest asset on an insurer's balance sheet. In today's Enterprise Risk Management environment, this credit risk needs to be managed better than it has often been in the past.

So far, most of this paper has been written from the perspective of the actuary involved in pricing (and therefore likely selling) reinsurance. Credit risks are primarily an issue when you are *buying* reinsurance, so this section is written largely from the buyer's perspective.

However, it can also go the other way too – for example, if your contract has adjustable premium features then you have credit risks from your cedant with regards to these future payments (albeit often this will be mitigated by an offset clause in the slip¹⁰). We wanted to include this topic in our paper as it can have implications for the pricing of reinsurance.

a – Pricing approaches

Theoretically, insurers could seek premiums from their reinsurers to compensate for the credit risk they are assuming by trading with them. In practice of course it could never be this blunt:

“Dear Mr Reinsurer, please pay us some money if you want us to buy your product, because we are concerned you might not be around in the long run.”

Hmm, short conversation, that one. However, it does effectively happen this way where cedants manage their credit risk by demanding collateralisation of the reinsurance limit from certain reinsurers, either partially or in full, often with an LOC. (This used to be widespread in the London Market, especially for UK companies reinsuring US cedants.) The reinsurer pays a fee to the bank for establishing the LOC so it is effectively paying a premium to cover its own credit risk, albeit to the bank rather than directly to the insurer.

Another theoretical possibility would be differentiating the premium rate that the insurer is willing to pay for each subscribed line to reflect the different credit risks. For example, the insurer might pay 100 to an AAA-rated reinsurer but only, say, 95 to a lesser rated reinsurer for the same cover. In practice this doesn't happen. (Incidentally, the differential pricing in the Aviation market is driven more by volumes of trade than by anything to do with creditworthiness.)

In practice P&C insurers usually take a simple view that above a specified rating (which might vary by class of business, but is often “A” for long tail risks) all reinsurers are good trading partners and should be treated equally. The insurers will refer to an in-house “approved security list” from which, in general, any reinsurer can be used with no explicit pricing adjustment for credit risk.

¹⁰ Did we say how important it is to read the slip yet? An offset clause will allow a reinsurer to offset monies it owes the cedant (i.e. recoveries) against monies they owe the reinsurer (i.e. premium adjustments).

Given the sums involved, the big issue really isn't about whether reinsurer ABC requires a 2% rate reduction, it's more about making sure we don't use ABC if we are concerned about their credit worthiness. Would you rather pay £20,000 for a car that might just let you down in the middle of the night in the pouring rain and freezing cold when you are miles from civilisation, or £20,500 for an identical car that has is guaranteed not to let you down? When the fractionally cheaper car breaks down will you think "*that's okay, this is why I saved £500*" or will you regret your decision to buy it in the first place? It's the same with reinsurance – only more so!

One common reinsurance feature that has a direct bearing on the credit risk is a "downgrade clause". This is a clause in the slip that will enact a change in the contract in the event that a certain credit rating falls below a particular threshold, for example if the insured loses its "A" rating from S&P. On property treaties such clauses typically cut off cover immediately, the insured gets a pro-rata refund of premium and the reinsurer caps off their exposure. With longer tailed liability classes this is unlikely to do much good – too little too late – so it is more usual to see the downgrade clause trigger collateralisation of the reinsurance limit. We said it before but we'll say it again – Read The Slip!

b – Considerations

Market events lead to capacity crunches. Just when you need your reinsurance most (i.e. when you have a large claim or series of claims), the reinsurers' finances are probably being stretched very hard;

There are other considerations that may affect an insurer's ability to pay besides the financial strength reflected on paper. It is difficult to quantify what effect these should have on the pricing as they depend on individual circumstance but they are worthy of thought:

- How exposed are they to volatile classes of business and/or concentrations of risks that might increase the likelihood of financial disaster?
- How reliant are they on retrocession, and who are they ceding large amounts to?
- Do they have a history of disputing claims? Some reinsurers gain a reputation for being the slower or faster payers and also for being more or less likely to dispute claims. Often becoming a slower payer or a more-likely-disputer is a sign of weakening finances...
- Who is the parent and what level of support and guarantees do they provide?
- A subtle point this one: In most of the US States and EU countries, insurance risks rank before reinsurance risks in the order in which insolvencies are paid. If your reinsurer writes a large book of insurance business you are carrying a greater credit risk than if they didn't (all other things being equal) because there is a large block of preferred creditors standing between you and your money. (This may only be a second-order effect in practice?)
- Interactions and correlations *between* reinsurers is an important dynamic. Most if not all of the major reinsurers have heavy exposures to the same risks, and the "retrocession web" is thick and tangled. Hurricane Katrina rocked a lot of people hard in 2005. So did Andrew in 1992. A quick succession of two or three massive events could bring down a number of reinsurers at once, and if a Very Big Reinsurer were to tumble this would no doubt create many further casualties in the reinsurance market as all their reinsurance customers saw their program significantly weakened.

A good paper was presented to GIRO in 1996 (“Reinsurance Security” by Sanders et al) which is worth a read. This sets out the important issues from the perspective of the reserving actuary – bad debt provisions and so on – but the parallels with pricing are clear.

We should also look to the banking community to see what we can learn to help us here.

We believe that this whole area of credit risks warrants specific and thoughtful actuarial research and development.

c – Modelling credit risks using CDS prices

Modelling credit risks well is difficult and we are not aware of any established robust method that works for reinsurance. This is an area worthy of exploration so we started to give it some thought, but as yet we have only really made a light scratch at one part of the surface!

The fair discount for credit risk versus a fully collateralised reinsurance price should be equal to the price of the respective credit hedge. Credit Default Swaps¹¹ are a readily available window to the market price for a credit hedge so how might they be employed to help us?

One approach could be to combine the estimated loss cost and payment pattern of the reinsurance with a market-based credit curve to calculate a market-consistent estimate of the value of the hedge:

Example (ignoring inflation)

Expected loss cost = £10m

Year	1	2	3	4	5	6	7	8	9	10
Cumulative Payment as % of Ultimate	0%	1%	5%	15%	50%	65%	75%	85%	95%	100%
Exposure £ remaining at start of Year	10m	9.9m	9.5m	8.5m	5m	3.5m	2.5m	1.5m	500k	-

Due to the annual nature of the CDS, to hedge these exposures we require the following CDS contracts:

- Term 9 years, Amount £500k
- Term 8 years, Amount £1m
- Term 7 years, Amount £1m
- Term 6 years, Amount £1m
- Term 5 years, Amount £1.5m
-

¹¹ CDS are basic instruments used to hedge credit exposure in investment banking. Basically, in return for an annual premium, a CDS pays out predefined amount if the issuer defaults on bond repayments within a given term. (A fuller definition of CDS contracts can be found in CA1.)

Suppose the market prices for these CDS contract terms on an A-rated company are as follows:

Year	1	2	3	4	5	6	7	8	9	10
Cost of CDS p.a. (A Rated)	0.08%	0.11%	0.14%	0.18%	0.21%	0.24%	0.27%	0.29%	0.31%	0.33%

Cross-multiplying these rates with the respective notional amounts for each term gives the total estimated market cost of the hedge. (This is of course an imperfect hedge as it does not cover the uncertainties in the amount or the timing of the notional amounts relative to our assumed pattern.)

In this example the total cost of hedging against an A-rated reinsurer works out as £116,000 or 1.6% of the expected loss cost. Obviously, for a reinsurer with a lesser rating you would have higher CDS prices and vice versa, and also the longer the payment pattern the greater the likely cost of the hedge and vice versa.

Of course, this method has a number of clear limitations:

- It assumes the RI recoveries are known in advance;
- It assumes the payment pattern of these recoveries is annual and known in advance;
- It ignores the likelihood of making a partial recovery in the event of insolvency;
- Using credit-default based rates as a proxy for reinsurance default is likely to underestimate the frequency of default as they don't really capture all of the risks to reinsurance default (e.g. won't pay as opposed to can't pay), but countering that CDS prices are "digital", i.e. they do not factor in the possibility of making a partial recovery upon default.

The approach could easily be performed stochastically in order to model your way around some of these issues, but we recognise it still isn't perfect.

It is important to take care when applying banking products to a reinsurance environment. Below we have outlined some issues (there are bound to be others) to be aware of in using Credit Default Swaps which are probably applicable to most investment banking credit solutions:

- The Credit Default Swap is not an insurance instrument and as such does not indemnify;
- Bond default rates do not allow for the 'won't pay' element, nor do they reflect the likelihood that a company will start disputing and non-paying its liabilities before it starts defaulting on its bond holders;
- The CDS may only cover the parent, so there is a risk of reinsurance default by a branch or subsidiary reinsurer that would not trigger the derivative, meaning that the CDS price could well underestimate the default probability for the subsidiary if there is not an absolute parental guarantee;

- The CDS is not triggered by the event of a company going into run-off. Reinsurers in run-off are often very slow payers who dispute more often than usual, which increases and extends the credit risk;
- The CDS contract still includes an element of credit risk against the issuer. However it may be unlikely that both the issuer and the reinsurer go bust at the same time.
- You might not be able to get a CDS price for your particular reinsurer, although you may be able to find a “similarly secure” reinsurer for whom you can.

Alternatives?

Alternative methodologies could possibly be developed around the idea of “capital required to support the business over the run-off period, allowing for default” looking to build on recent papers by Don Mango?

As noted before, we believe credit risks warrant further study.

10) SO WHERE DO WE GO FROM HERE?

There is a limit to the amount that even the most diligent group can fit into a single working party “season”, and it is by no means unusual to find at the end that there is scope to expand upon and develop what has been done, given more time. We find ourselves in this position. Our journey (including the preparation of our “reading guide” referred to at the beginning) revealed a number of areas where we believe there is scope for valuable research and development in the non-life actuarial arena. In particular we would suggest the following:

- We hope that our reading list can be incorporated into the Institute’s official reading guide and circulated to raise awareness of just how much “good stuff” there is readily available on the subject of reinsurance. We also feel that this resource should be regularly updated as new papers are written and valuable papers are noted as being missing from our list.
- We recognise that paper documents can rapidly become obsolete, so perhaps the Institute might consider developing “wiki” style websites where practitioners can share knowledge and experience in an interactive manner? For example, this pricing paper could form the starting basis for a “reinsurance pricing issues” wiki and develop into a living, evolving resource.
- We did try to gain access to Lloyd’s market data in order to build some market-level benchmarks, such as ILFs for a few key classes. Lloyd’s were initially receptive but it turned out to be impractical to provide the data that we would have needed. This is a shame, and we shouldn’t just forget the idea. Should the Institute lend its might to sponsoring an initiative along these lines, using data from the ABI or from Lloyd’s or from Xchanging or wherever makes sense? Our American counterparts have a lot of market curves and factors available to them, largely thanks to the ISO (Insurance Services Office). Why shouldn’t we?
- Correlations, accumulations and aggregations. This is an important area from many perspectives: reinsurance pricing of course, but also reinsurance design, reserving, capital management, risk management and so on. We recognise that there are many technically minded actuaries who are expert in these things but there are also a great many less technical actuaries who are not. An “everything you *really* need to know about working with correlations and accumulations *in the real world*” paper would be a fantastic resource, wouldn’t it? Maybe this could be a part of a bigger scope looking at stochastic modelling?
- Credit risk is another area fundamental to our world but which is too often skimmed over. The typical “street actuary” (ourselves included) could benefit from boosting this area of their toolkit. The banking world and our actuarial colleagues in the investment field are but two sources begging to be tapped. Again, an “everything you really need...” type paper would seem to us to be a very worthwhile pursuit.

And there we have it, the end, just one appendix left. Thanks for reading our paper – we sincerely hope you found it enjoyable and worthwhile!

The 2006 GIRO “Reinsurance Matters!” Working Party

Mark Flower (Chairman); Ian Cook; Craig Divitt; Visesh Gosrani; Andrew Gray; Gillian James;
Gurpreet Johal; Mark Julian; Lawrence Lee; David Maneval; Roger Massey

APPENDIX

A LITTLE MORE ON CREDIT RISKS...

Default rates

So why should all reinsurers above a certain rating be classed as a similar risk?

In January 2000 an Institute of Actuaries working party produced a very good paper entitled “Reinsurance Bad Debt Provisions for General Insurance Companies”. This was subsequently updated in October 2005 and is now effectively an Advisory Note, having been incorporated as part of GN18. The paper includes tables of default rates for bonds of different ratings and shows how the ‘Static Pool Average Cumulative Default Rates’ calculated by S&P, go up for each rating as time increases.

These bond default rates are often used as a proxy for the probability of a reinsurer with a given rating defaulting on payment of a valid claim for reserving purposes.

Old Table:	Static Pool Cumulative Average Default Rates (%)					(Data to 1997)				
	Yr1	Yr2	Yr3	Yr4	Yr5	Yr6	Yr7	Yr8	Yr9	Yr10
AAA	0	0	0.06	0.12	0.19	0.35	0.52	0.82	0.93	1.06
AA	0	0.02	0.1	0.2	0.35	0.53	0.7	0.84	0.91	1
A	0.05	0.14	0.24	0.39	0.58	0.77	0.97	1.22	1.49	1.76
BBB	0.18	0.42	0.67	1.21	1.68	2.18	2.66	3.07	3.38	3.71
BB	0.91	2.95	5.15	7.32	9.25	11.22	12.29	13.4	14.33	15.07
B	4.74	9.91	14.29	17.42	19.7	21.26	22.56	23.75	24.71	25.55
CCC	18.9	26.01	30.99	35.1	39.02	39.88	40.87	41.17	41.86	42.72

When the paper was originally published in 2000 the cumulative default rates were based on data up to 1997 (see table above). What this table does show is that there is a big difference in default rates for bonds rated BB or worse compared to those with BBB or better ratings (so called “investment grade”), so would it be unreasonable to group together all those companies with a rating of, say, A or better and treat them as a similar credit risk? The default rates for bonds rated A and better are pretty small and the variance between the average default rate of an A rated bond and a AAA rated is less than a factor of 2, so perhaps it is reasonable to consider them all equally “good” when pricing reinsurance, for example?

When the Advisory Note was updated in October 2005 a new table of cumulative default rates was included which took account of experience up to 2004. There are some quite large differences between the old and new default rates provided (see tables below).

Old Table:	Static Pool Cumulative Average Default Rates (%)										(Data to 1997)
	Yr1	Yr2	Yr3	yr4	Yr5	Yr6	Yr7	Yr8	Yr9	Yr10	
AAA	0	0	0.06	0.12	0.19	0.35	0.52	0.82	0.93	1.06	
AA	0	0.02	0.1	0.2	0.35	0.53	0.7	0.84	0.91	1	
A	0.05	0.14	0.24	0.39	0.58	0.77	0.97	1.22	1.49	1.76	
BBB	0.18	0.42	0.67	1.21	1.68	2.18	2.66	3.07	3.38	3.71	
BB	0.91	2.95	5.15	7.32	9.25	11.22	12.29	13.4	14.33	15.07	
B	4.74	9.91	14.29	17.42	19.7	21.26	22.56	23.75	24.71	25.55	
CCC	18.9	26.01	30.99	35.1	39.02	39.88	40.87	41.17	41.86	42.72	

New Table:	Static Pool Cumulative Average Default Rates (%)										(Data 1981 to 2004)
	Yr1	Yr2	Yr3	yr4	Yr5	Yr6	Yr7	Yr8	Yr9	Yr10	
AAA	0	0	0.04	0.07	0.12	0.21	0.31	0.48	0.54	0.62	
AA	0.01	0.03	0.08	0.16	0.26	0.4	0.56	0.71	0.83	0.97	
A	0.04	0.13	0.26	0.43	0.66	0.9	1.16	1.41	1.71	2.01	
BBB	0.29	0.86	1.48	2.37	3.25	4.15	4.88	5.6	6.21	6.95	
BB	1.28	3.96	7.32	10.51	13.36	16.32	18.84	21.11	23.22	24.84	
B	6.24	14.33	21.57	27.47	31.87	35.47	38.71	41.69	43.92	46.27	
CCC	32.35	42.35	48.66	53.65	59.49	62.19	63.37	64.1	67.78	70.8	

	% Increase in Average Default Rate From Old Table to New Table										
	Yr1	Yr2	Yr3	yr4	Yr5	Yr6	Yr7	Yr8	Yr9	Yr10	
AAA			-33%	-42%	-37%	-40%	-40%	-41%	-42%	-42%	
AA		50%	-20%	-20%	-26%	-25%	-20%	-15%	-9%	-3%	
A	-20%	-7%	8%	10%	14%	17%	20%	16%	15%	14%	
BBB	61%	105%	121%	96%	93%	90%	83%	82%	84%	87%	
BB	41%	34%	42%	44%	44%	45%	53%	58%	62%	65%	
B	32%	45%	51%	58%	62%	67%	72%	76%	78%	81%	
CCC	71%	63%	57%	53%	52%	56%	55%	56%	62%	66%	

By comparing the new rates with the old it is apparent that the default rates for AAA rated bonds have decreased significantly whereas the default rate for A rated bonds has increased. (However, note that the number of AAA-rated reinsurers has also reduced significantly!)

If we are to assume that bond default rates are a good proxy for reinsurers defaulting on valid claims then the ‘new’ data would suggest that there is now less justification for treating all reinsurers with a rating of A or better as the same credit risk and perhaps cedants should accept, and be prepared to pay a premium for AAA reinsurance cover?

The other interesting change in the default rates is the substantial increase in default rates for bonds rated BBB or worse. The 5-year average cumulative default rates have typically increased by more than 50%. Again, if we are to assume that bond default rates are a good proxy for reinsurers defaulting on valid claims then the change between the ‘new’ and ‘old’ default rates would suggest that credit risk for reinsurers with lower ratings has increased a lot over the fairly short period from 1997 to 2004. Since the bulk of reinsurance bad debt provisions are in respect of reinsurers with ratings of BBB or worse, if the default rates are used blindly in calculating this provision it would suggest that purely due to the change in default rate tables the provision could be expected to increase by in excess of 50%.

If using the impairment rates observed by rating agencies to put a value on the credit risk on a contract, a number of factors should be considered, including:

- The rates are for the combined property/casualty and life/health sector
- Financial impairment as defined by rating agencies is different to a measure of insurers that would default on policyholder obligations.

Finally, we would observe that the increases in default rates should not be surprising when you consider the number of re/insurer insolvencies we saw in the period from 1997 to 2004!