

P. Parodi – Pricing in general insurance

Solutions to selected problems

This is a collection of solutions to selected problems from the textbook. The problems included in this collection tend to be those that are more complex and that make for more suitable as examination material, although hints for other problems are sometimes offered.

More solutions will be added over time.

Please contact the author (pietro_parodi@yahoo.com) if you wish to report errors in the solutions.

Chapter 1

Question 1

Assumptions:

- Claims inflation: 4%
- Data as at: 12/6/2011
- No IBNER
- No IBNR except for 2010
- ... it's property anyway so IBNR shouldn't normally be a significant concern
- From 2008 onwards, a new smoke detection system has been installed which is thought to decrease the probability of a fire by 20%
- Number of breweries is constant (and equal to the current number, three)
- Loading for expenses: 10%
- Loading for profit: 10%

First of all we need to align the exposure data (which are on a fiscal year basis) to the claims data (on a policy year basis)

Fiscal year (1/4)	Number of pubs	No of pubs (policy year)
2006/7	653	682.0
2007/8	711	704.5
2008/9	698	700.0
2009/10	702	706.0
2010/11	710	715.0
2011/12	720	720.0

Then we need to correct the claims data figure by inflation, by exposure and by risk profile changes, to bring it in line with the renewal policy year (Columns 6 and 7 below).

Policy year (1/10)	Total paid	Total O/S	Recoveries	aid+O/S-Recoveries	Policy year	Inflated to 1/4/2012	Exposure-adjusted	Risk-adjusted (new alarm system)
2006/2007	3,266,198	-	97,381	3,168,817	2006	3,855,351	4,070,165	3,256,132
2007/2008	3,148,272	4,617	728	3,152,162	2007	3,687,583	3,768,716	3,014,972
2008/2009	1,608,092	21,182	1,987	1,627,287	2008	1,830,476	1,882,775	1,882,775
2009/2010	1,324,927	589,309	2,619	1,911,617	2009	2,067,605	2,108,606	2,108,606
2010/2011	93,203	975,322	-	1,068,524	2010	1,111,265	1,119,036	1,119,036
Grand Total	9,440,692	1,590,430	102,715	10,928,407	<i>Average (2006-10)</i>	2,510,456	2,589,860	2,276,304
					<i>Average (2006-09)</i>	2,860,254	2,957,565	2,565,621
					<i>Loadings (10% exp, 10% profits)</i>		3,696,957	3,207,027

Finally, we assume that the introduction of the fire alarm system in 2008 leads to a reduction of 20% in the frequency and hence in the overall amount of claims. For this reason, we correct the claim amount for 2006 and 2007 so that it is equal to what the experience would have been if the fire alarm had already been in place there.

We exclude year 2010 because it will have some IBNR

We calculate the average over the remaining years, obtaining Average = GBP 2.57m.

We apply loadings for expenses and profits obtaining

Tech premium = $\text{GBP } 2.57\text{m} / (1 - 0.2) = \text{GBP } 3.17\text{m}$

(a tech premium calculated as $\text{GBP } 2.57\text{m} / (1 - 0.2) = \text{GBP } 3.21\text{m}$ as in the table above would have been equally acceptable)

(b)

- The analysis is quite crude as it is based on a simple burning cost exercise over 5 years!
- All the losses seem relatively small. We have made no adjustment for the possibility of large losses.
- Experience rating seems a particularly poor method here because of the lack of large losses and the very concrete possibilities that most if not all the losses come from pubs – hence, the presence of the breweries may not be priced in at all!
- If you have given full credence to the 20% reduction in probability, you have probably believed too readily the message that your client (or his broker) has given to you – the impact might be well lower. You should state that.
- If you have given it no credence, you might have been too pessimistic and in any case you should state that assumption

Chapter 2

Most of the questions in this chapter can be answered by referring to the text or by looking up the information on the web.

Question 2 (part)

Both employers' liability and workers' compensation have the ultimate social purpose of protecting employees against the consequences of accidents and diseases in the workplace.

The victim will be compensated through an employers' liability policy if the employer has been negligent...

... whereas workers' compensation is completed unrelated from questions of negligence – it takes care of the consequences of accidents/diseases without a need to prove negligence.

In practice, the distinction between employers' liability and workers' compensation is not so sharp in the way they're applied, since the definition of negligence has become so broad as to be very difficult for an employer not to be legally liable for virtually anything that happens accidentally to its employees.

There is also a territorial difference: territories tend to have either workers' comp (as in the US) or employers' liability (as in the UK)

Question 10

Hint: It is the insured who is restored to the former financial state, not the victim!

Question 11

The legal framework is that of tort law.

There are many categories of tort, but the most important for liability insurance is that of *negligence*...

...i.e. the idea that someone had a duty of care towards their neighbour and has somehow failed to exercise it.

Liability insurance indemnifies the insured for legal liability related to bodily injury or property damage caused by the insured to a third party because of negligence.

Chapter 3

Question 1

Motor

A popular reinsurance purchase for motor insurance in the UK is excess of loss insurance, although other arrangements such as quota share are possible. This is because the main risk in motor insurance is that of very large liability claims (payouts of up to £40m have been seen to happen in motor) Aggregate/Cat XL may also be purchased to protect against large-scale catastrophe events (most notably flood in the UK)

Property

Typical products purchased are

- Excess of loss insurance
- Aggregate/Cat XL
- Surplus reinsurance

Again, the aim is to protect against very large individual losses such as a fire at a very expensive plant (XoL), and to protect against the effect of catastrophes affecting a large number of properties.

Surplus offers an interesting way for the company to mitigate the larger risks in its portfolio while keeping most of the smaller risks, by choosing how much to retain for each risk

Liability

- XL reinsurance
- Aggregate XL

XL reinsurance is purchased for the same reasons as for motor and property.

Aggregate XL offers the possibility of putting together claims arising from the same cause but not necessarily with large individual amounts.

Combined book

For the combined account, the company may consider stop loss reinsurance (if it is able to purchase it), which protects the overall account (or a subset of it) against the possibility of a high loss ratio and therefore further stabilises the balance sheet.

Other considerations

Medium-sized, so probably no need for quota share

Industry loss warranty (protection based on an industry index rather than on the company's performance) is also a possibility

Question 2

(a)

In treaty reinsurance, a LOD policy is one where the reinsured is covered for all claims occurring during a given period, without regard to the inception date of the original policy and the reporting date

In treaty reinsurance, a RAD policy is one where the reinsured is covered for all claims related to the original policies incepting over a given period, without regard to the loss date and the reporting date.

In treaty reinsurance, a claims-made policy is one where the reinsured is covered for all claims reported over a given period, without regard to the loss date and the inception date of the original policy

(b)

One way to carry out the calculation is as follows:

- (a) First of all, we need to check that the claim is a valid one for the original policy. Some of the claims in the list (possibly because of some error in the original statement of the problem!) have both a loss date *and* a reporting date prior to the inception of the original policy, which is not admissible (assuming the original policy was either on an occurrence or a claims-made basis)
- (b) Secondly, the amount which is relevant to the purchased layer is calculated by the formula $Amt_to_the_layer = \min(L, \max(X-D, 0))$ where $X = \text{loss}$, $D = \text{£}0.5\text{m}$, $L = \text{£}1.5\text{m}$
- (c) Thirdly, for each basis we only consider the losses that fall within the time period 1/10/2012 – 30/9/2013 with regard to the original policy inception date (RAD), the loss date (LOD) and the reporting date (CM)

The following table shows the detailed calculations:

Claim ID	Inception date of original policy		Reporting date	Amount [£]	Policy year (RAD)	Policy year (LOD)	Policy year (CM)	Valid claim in the Include			Loss to the layer [£]	
	Loss date	original policy						original policy	d in RAD	Included in LOD		Included in CM
1	01/07/2012	24/03/2013	08/12/2013	647,935	2011	2012	2013	1	-	1	-	147,935
2	01/10/2012	21/02/2013	18/08/2013	204,973	2012	2012	2012	1	1	1	1	-
3	01/10/2012	03/07/2012	03/07/2012	179,985	2012	2011	2011	0	-	-	-	-
4	01/01/2012	19/12/2012	17/03/2013	812,868	2011	2012	2012	1	-	1	1	312,868
5	01/01/2012	11/02/2013	09/10/2013	37,036	2011	2012	2013	1	-	1	-	-
6	01/07/2012	28/01/2012	03/02/2012	2,058,075	2011	2011	2011	0	-	-	-	-
7	01/10/2012	14/01/2013	13/08/2013	219,190	2012	2012	2012	1	1	1	1	-
8	01/01/2012	20/02/2013	06/03/2013	665,041	2011	2012	2012	1	-	1	1	165,041
9	01/04/2012	11/08/2013	07/02/2014	1,152,750	2011	2012	2013	1	-	1	-	652,750
Total amount [£]								-	1,278,594	477,910		
								(RAD)	(LOD)	(CM)		

(c)

The exact timing of the payments is irrelevant – only the settlement time matters for the LMIC

The index at the time of settlement is 157, therefore the deductible is going to be $D' = D \times I(\text{settl})/I(\text{inception}) = \text{£}500,000 \times 157/143 = \text{£}548,951$.

The value of the limit doesn't matter in this case because it is already above the loss amount. In any case, it's three times the amount above, i.e. £1,646,853

The amount ceded to the layer is therefore $\min(L', \max(0, X-D')) = \max(0, X-D') = \max(0, 1152750 - 548951) = \text{£}603,799$

Chapter 4

Question 8

(a)

Product liability.	Probably is crucial as it covers against the adverse effect of drugs, which may lead to very large claims
Public liability	The company needs to protect itself against generic injury/damages against third parties because of anything not covered by other policies. It may be rolled into one with product liability.
Employer's Liability	This is a compulsory cover that indemnifies the insured against legal liability towards its employees
Environmental liability	The company might eg pollute rivers by offloading chemical waste used for drug production
Motor insurance – Fleet and commercial vehicles	The company will certainly have a fleet of car (e.g. pharma representatives, directors, etc.) and commercial vehicles for e.g. the transportation of materials and for distribution
Commercial property...	The company will need to protect plants, warehouses and offices from various hazards such as fires, floods and other man-made and natural catastrophes
... and business interruption (aka consequential loss)	If there is a damage to the company's property, this may have the collateral effect of stopping production
Directors and officers	The company's director may be held accountable for various wrongful acts, e.g. publication of financial statements
Marine – Goods in transit	The company might need to protect against loss, damage or theft of transported pharmaceutical products

(b)

First of all, notice that none of these products are compulsory except for employer's liability and third party liability for vehicles.

Among the non-compulsory products, the overriding consideration is that we have a large UK pharmaceutical company, which may well think it has the financial strength to retain its own losses

In particular, the company may think they'll have no difficulty dealing with small, predictable losses, such as own-damage to their fleet

The insurance market might be reluctant in a given period to insure products which may result in very large payouts, especially product liability

The price for a given product may be too large, pushing the company towards self-retention

Even if the price of a product is reasonable, the company's perception may be that they can invest in risk control mechanisms to push down the expected costs of a given risk (e.g., incentive schemes for personnel, or good risk prevention measures against fire)

Note that although all these answers refer to a binary decision (buy/not buy) many of them apply to the question of whether we should retain losses below a certain level, or some other retention structure.

(c)

See Section 4.1.2.3

Chapter 5

Most of these questions can be answered by doing some research on the web.

Question2

E.g.: earthquake and hurricanes for the US, windstorm and flood for the UK, earthquake and flood for Italy...

Question3

Space weather tends to be more predictable since most of it has to do with solar activity which is subject to a 11-year cycle (22 years if you take the direction of the magnetic field in the sun into account). When the number of sunspots increases, the disruptive effects of solar weather are felt more keenly.

There are, of course, other events such as the explosion of supernovae that could cause damage and that are not related to the solar cycle – this reduces the predictability of space weather.

Chapter 6**Question 1**

As recalled in the text, the various moments of S can be expressed as different derivative of $M_S(t)$ calculated at $t = 0$. Using Equation (26.5) and the properties of the derivative of compound functions (in this case $\left(f(g(h(x)))\right)' = f'(g(h(x))) \times g'(h(x)) \times h'(x)$), we can write:

$$E(S) = M'_S(t)|_{t=0} = M'_N(\log M_X(0)) \times \frac{1}{M_X(0)} \times M_X^{(0)} = E(X) \times 1 \times E(N)$$

And similarly with $Var(S)$ after writing it as $Var(S) = E(S^2) - E(S)^2 = M''_S(0) - M'_S(0)^2 = \dots$.

(This is a rather dry exercise with little practical use but it helps the understanding of what S is and helps remember the formula for $E(S)$ and $Var(S)$.)

Question 2

(a)

The collective risk model is a model for the total losses that an entity incurs over a given period.

According to this model, the total losses incurred during a period T is a random variable S that can be written as $S = X(1) + \dots + X(N)$ where N is the number of losses (itself a random variable) and $X(k)$ is the amount (again, a random variable) of the k -th loss.

The main assumptions of the CRM are that:

- The loss amounts $X(1), X(2), \dots, X(N)$ are independent of the number of losses
 - The loss amounts $X(1), \dots, X(N)$ are independent and identically distributed
-

(b)

S represents in this case the total losses incurred by the company because of their liability towards their employees.

The assumption that the losses are identically distributed does not present difficulties – if we accept that the losses are the result of some stochastic process, there will always be a unique severity distribution associated with it. Of course different types of losses will ultimately be the product of different subprocesses, but the severity distribution that comes from mixing the severity distributions of the different subprocesses will still be unique for the combined process.

The assumption that the losses are independent of each other is likely to be true only approximately: e.g. if many of these liability losses came from related events or from an underlying problem, they would certainly be somehow correlated

The assumption that N is independent of $X(1), \dots, X(N)$ is also an approximation, and it is easy to imagine situations in which this assumption breaks down. Eg, if a large percentage of the company's employees have to work outdoor, poor weather conditions in a given year may mean more injuries in outdoor activities but not more accidents in indoor activities, and since outdoor and indoor activities will in general have different severity distributions, the overall severity will change as a consequence of the larger number of losses.

(c)

$$E(S) = E(X)E(N)$$

$$\text{Var}(S) = \text{Var}(X)E(N) + (E(X))^2 \text{Var}(N)$$

(d)

MC simulation works by generating a large number, say 100,000, of simulated years of losses

For each year k , we draw a random from the frequency distribution, eg a Poisson or Negative binomial

For each of the $N(k)$ losses in a particular simulated year, we draw a random number from the severity distribution, obtaining $N(k)$ simulated losses: $X(1), \dots, X(N(k))$

Then we sum all the losses thus sampled, obtaining the total losses for that simulated year: $S(k) = X(1) + \dots + X(N(k))$

Finally we sort all the $S(k)$ in ascending order, thus obtaining an empirical total loss distribution $F_{emp}(s)$

The 95th percentile can be calculated by picking the 95th percentile in such a sorted list, e.g. if there are 100,000 elements, the 95th percentile will be roughly element 95,000.

(e)

Other common methods that people use to produce an aggregate loss distribution are the Gaussian approximation, the translated Gamma approximation, Panjer recursion and Fast Fourier transform.

Question 3

This is mostly bookwork based on Section 6.3 – but it helps trying to write down the flowchart without referring to the book.

Chapter 8

Rather broad questions that can be answered by referring to various chapters in the text and that should help readers develop their thinking about the role of data in pricing.

Question 2

Outline of a solution

Apart from the standard information (see, e.g., Figure 8.3), we would ideally include:

- information on the type of plane (e.g. jet vs turbo prop)
- airliner and territory of airliner
- territory of crash
- passengers profile (nationality, etc)
- ...

All this information will help establish which losses are relevant when pricing aviation for a given airliner, and will provide rating factors.

A lot of this information is publicly available since large aviation losses attract plenty of media attention

Chapter 10

Question 4

a.

Claims inflation is that phenomenon by which a claim which happened, say, two years ago, would have a different value were it to occur today because of inflationary/deflationary pressures of different types.

These inflationary pressures will be different, in general, for different types of claims. E.g. for motor claims, the heads of damage may include such things as repair costs, car replacement values, loss of earning, medical care, etc. which do not necessarily grow in line with consumer price inflation

b.

Motor bodily injury claims have different heads of damage. The most important are:

- medical care → related to inflation of cost of care (specialised indices)
- losses of earning → related to wage inflation (AWE indices in the UK)

These may not account for the full value of claims inflation, because, e.g., of “court inflation”: courts show a tendency to increase compensation over the years over and above the effect of the other costs above

Furthermore, large claims normally are subject to larger inflation and this may push up the overall inflation from the ground up

c.

First, each claim must be inflated to the mid-term of the policy:

Claim amount revalued (FRD) = claim amount (FRD) x Index (@ mid-term of policy) / Index (@ loss date)

Then, all claims must be converted to GBP using today’s currency exchange rate (1 FRD = 0.0164 GBP)

Actual calculations:

Claim_1 (revalued, FRD) = FRD 99,799 x 213 / 102 = FRD 208,403
 Claim_1 (revalued, GBP) = FRD 208,403 x 0.0164 GBP/FRD = GBP 3,418

All calculations:

<i>Claim ID</i>	<i>Loss date</i>	<i>Original currency</i>	<i>Claim amount [FRD]</i>	<i>Claims inflation index at loss date</i>	<i>Exchange rate at loss date (value of 1 FRD in GBP)</i>	<i>Claim amount revalued (FRD)</i>	<i>Claim amount revalued (GBP)</i>
Claim_1	11/01/2008	FRD	99,799	102	0.0891	208,403	3,418
Claim_2	03/03/2009	FRD	2,665,226	117	0.0724	4,852,078	79,574
Claim_3	03/10/2009	FRD	510,466	133	0.0581	817,513	13,407
Claim_4	15/06/2011	FRD	432,408	180	0.0222	511,683	8,392
Claim_5	21/07/2011	FRD	921,064	186	0.0191	1,054,767	17,298
	<i>Today</i>			<i>201</i>	<i>0.0164</i>		
	<i>Mid-term of policy</i>		<i>(projected)</i>	<i>213</i>			

d.

The purpose of this combined claims revaluation and conversion is to estimate what value (in GBP) a claim which has occurred in the past (in FRD) would have if it were to happen at the mid-term of the policy.

Now if claim_1 were to happen in the mid-term of the policy, first of all it would be revalued to FRD 208,403 because that claim would still happen in Freedonia!

Then, it must be converted into the GBP amount that it will have at mid-term of policy, by applying the exchange rate at the mid-term of the policy

However, we cannot know the future exchange rate, and since this is a traded quantity, the current exchange rate of today should theoretically incorporate all information about that currency...

... and therefore today's ROE is the best available estimate for the ROE at the mid-term of the policy

Note that the currency exchange rate at the time of the loss is not relevant

Chapter 11

Question 4

(a)

EEL is each and every loss deductible. It works like this: if the insured incurs a loss of X, the first EEL of this loss is retained by the insured, and the rest is ceded to the insurer.

Therefore:

- Gross loss = X
- Retained = $\min(X, \text{EEL})$
- Ceded = $\max(X - \text{EEL}, 0)$

For example, if $X = \text{£}150,000$, the first $\text{£}100,000$ are retained and the remaining $\text{£}50,000$ are ceded.

Having an EEL allows the insured to retain the small losses and focus on the big ones that can cause instability...

... and also keeps the costs of insurance down.

The annual aggregate deductible (AAD) modifies the policy so that the amount retained in a year never goes above a given amount.

As an example, if an EEL of $\text{£}100\text{k}$ is in place, with an AAD of $\text{£}250\text{k}$, if we have three losses

$X_1 = \text{£}150\text{k}$

$X_2 = \text{£}300\text{k}$

$X_3 = \text{£}80\text{k}$

The insured will retain $\text{£}100\text{k}$, $\text{£}100\text{k}$, and $\text{£}80\text{k}$ because of the EEL arrangement, but as the total amount is $\text{£}280\text{k} > \text{£}250\text{k}$, it will only end up retaining $\text{£}250\text{k}$ and the rest ($\text{£}30\text{k}$) will be ceded. The ceded amount can be calculated as $X_1 + X_2 + X_3 - \text{Retained} = \text{£}530\text{k} - \text{£}250\text{k} = \text{£}280\text{k}$.

The point of AAD is obviously to give more certainty of results to the insured.

The individual/aggregate limit is straightforward: whenever the total amount ceded to the insurer either through a single loss or the cumulation of different losses is above a certain level, say $\text{£}5\text{m}$, then the rest (after the retention bits have been taken out) will go back to the insured.

The main reason why a limit is in place is because the insurer wants to have some certainty around the largest amount that it may have to pay...

...however, it may also be in the interest of the client to reduce the limit so as to pay less premium...

... however, the insured must be really careful that the limit is fully adequate because breaking the limit may mean several millions of costs!

(b)

The claims are already revalued so we don't need to perform revaluation.

We have to determine how the policy is affected by these claims. To do that, we consider the retained losses and the ceded losses. To determine the retained losses for each year, first cap every individual loss to $\text{£}100\text{k}$ (by the formula $\text{capped}(x) = \min(\text{£}100\text{k}, x)$) then cap the total retained using the AAD.

Retained

	<u>2007</u>	<u>2008</u>	<u>2009</u>	<u>2010</u>	<u>2011</u>	<u>2012</u>
	64,995	56,284	7,115	2,237	4,801	97,938
	100,000	100,000	14,952	2,060	100,000	85,959
	1,010	27,036	8,576	25,106	55,957	
	100,000	2,559	100,000		71,563	
	33,277		100,000		75,680	
	30,319		6,582		2,634	
	3,297		59,881			
	100,000		536			
<i>Total</i>	432,897	185,879	297,643	29,403	310,635	183,897

(Alternatively, explain why this is not necessary to calculate the burning cost: the AAD will never be hit, since even the gross amount never goes up to £2m.)

The ceded losses are simply the total losses minus the retained losses.

(Alternatively, again, explain why this is not necessary to calculate the burning cost.)

Ceded

	<u>2005</u>	<u>2006</u>	<u>2007</u>	<u>2008</u>	<u>2009</u>	<u>2010</u>	<u>2011</u>
	-	-	-	-	-	-	-
	19,630	1,993	-	-	59,535	-	-
	-	-	-	-	-	-	-
	159,414	-	645,954	-	-	-	-
	-	-	246,021	-	-	-	-
	-	-	-	-	-	-	-
	-	-	-	-	-	-	-
	4,119	-	-	-	-	-	-
<i>Total</i>	183,162	1,993	891,975	-	59,535	-	-

We can use the totals above to estimate the expected losses for 2011. However, we have to make some adjustments.

First of all, we need to adjust for changes in exposure. However, in this case the historical exposure data are unavailable and we are told that the exposure has remained roughly unchanged, so we can assume that the exposure has remained constant over the loss experience period (and that it will remain unaltered also in the renewal period) and avoid any adjustments.

Let us calculate the average of the ceded losses

Accident year	Ceded
2007	183,162
2008	1,993
2009	891,975
2010	-
2011	59,535
2012	-
<i>Average 2007-11</i> [▼]	227,333
<i>Average 2007-12</i>	189,444

Some other considerations need to be made before calculating the technical premium:

- We have to assume that the risk profile of the company hasn't changed, e.g. that the proportion between clerical and manual workers has not changed. We have no information about whether that is the case, but we need to be explicit about that.
- We need to take IBNR somehow into consideration. There is not a huge IBNR here (most losses reported within a year, we're told – perhaps optimistically), but even if we buy the assumptions in the text we see that 2012 is probably very immature (the year hasn't even finished yet as the data are at 30/9/2012), and 2011 is technically not fully complete yet (losses between 1/10/2011 and 31/12/2011 may have not been reported yet).
- Based on the very limited information, *an acceptable approximation* is to assume that 2011 is roughly complete and to ignore 2012 altogether in the calculations.

The fact that this is only an approximation should be noted!

- We are told that some losses are only estimates. There might be some IBNER. However we have no information on this. We need to assume that IBNER = 0, but we need to be explicit about that.

(c)

	<i>Low</i>	<i>High</i>
<i>Average 2007-11</i>	227,333	227,333
<i>Target loss ratio</i>	70%	90%
<i>Technical premium</i>	324,761	252,592

Table 2. Calculation of the technical premium with target loss ratio at the extremes of the range.

Since we have no information on the payment/settlement pattern of claims, we have ignored investment income. Since risk-free return rates are currently quite low, this assumption does not have a big impact anyway.

The actual premium paid (£300k) is within the range of estimates in Table 1, and therefore appears in line with market practice and not extortionate. This is even more true given the uncertainties explained in (b).

This is even truer given the fact that often the burning cost analysis tends to underestimate the true cost.

The client might still want to walk away from that arrangement for other reasons – however, since this is EL (a compulsory cover in the UK) the scope for manoeuvre is limited.

(d)

There are 33 losses of which 7 above £100k – the EEL is getting rid of around 80% of the losses, which seems to be adequate (there is no hard and fast rule here, but the EEL is there to remove the very large majority of claims – 80% to 99% are common benchmarks).

The AAD is definitely too high considering that in order to be hit, we would need to have at least 20 claims above £100k in a year (or, say, 40 claims of average value £50k): the likelihood of something like this happening is extremely low.

To be useful, the AAD should be hit occasionally but not too often – say 1 in 5 years to 1 in 20 years. Looking at the total amount retained in the years 2007 to 2011, we see that £500k is probably a more adequate figure: it has never been hit but in year 2007 WWW got close to it.

As to the policy limit, £5m is the minimum amount of cover that an employer must have by law...

... however, losses of an average of £10m are not unheard of (*the text says there are two of them each year on average*) and if one of those were to hit WWW it may be felt on its balance sheet (*we do not have an exact idea of the company's profits but the turnover is not huge*).

Furthermore, most insurance companies offer EL with a limit of £10m anyway. It would therefore make sense to purchase at least £10m of cover.

As to the impact of these changes on the policy price:

- Decreasing the AAD normally implies an increase in the policy price, as a smaller amount is retained by the insured...
-
- ... however, since the proposed £500k does not have an impact on the burning cost, WWW might get away with a modest increase
-
- Increasing the limit should more surely bring a bigger price – using our burning cost model with the catastrophe assumption as stated for example brings an upward correction of around £50k (a loss of £10m will now have full impact instead of being limited to £5m)...
-
- ... however, if a more sophisticated model is used, the impact wibly be more nuanced.
-

Chapter 14

Question 4

(a)

Assumptions

- IBNR figures are fully relevant to the organisation, despite the fact that are on a portfolio level
- All claims reported as non-zero will stay as non-zero
- No other risk profile changes except for the change in exposure

Calculations

- To correct for IBNR, multiply the number of claim for each year by the portfolio development factor...
- The results are as in Column 1 below
- Then we can calculate the exposure-adjusted number of claims for each year by multiplying the number of claims for each year by the estimated 2015 exposure and divide by the historical exposure for that year
- The results are as in Column 2 below
- The default model for the number of claims where the collective model is applicable is the negative binomial (see textbook, Box 14.3).
- The parameters can be calculated with the method of moments. The mean is:

Mean = Straight average of exposure-adjusted, IBNR-corrected values
 (or Mean = Weighted average of IBNR-corrected values)
 = 219.5

- The variance with the method of moments is equal to the straightforward variance of the exposure-adjusted, IBNR-corrected values: $V = 858.9$
- The V/M ratio is 3.9
- The suggested model is therefore a negative binomial with $M=219.5$ and V/M ratio = 3.9
- An alternative representation of the NB in terms of p and r is welcome but not required.

Case 1

Ultimate	Ultim @ 2015
174.0	221.0
182.0	214.9
211.2	197.9
227.1	224.7
188.0	191.9
262.1	272.9
262.7	247.6
188.8	186.5
M	219.5
V	858.9
V/M ratio	3.9

(b)

- We correct for IBNR as in Column 1
- And for exposure as in Column 2
- Mean = 233.2
- Variance/mean ratio = 0.6
- Since these parameters are not consistent with a negative binomial model, we (following the textbook, Box 14.3) propose a Poisson model with rate = 233.2

Case 2

Ultimate	Ultim @ 2015
183.0	232.4
202.0	238.5
221.2	207.3
227.1	224.7
240.2	245.0
229.6	239.1
257.7	242.8
240.6	237.7
	233.2
	151.4
	0.6

Chapter 15

Question 1

(i)

Here's some of the methods that people use in practice or have been proposed
Ignore IBNER altogether, and use all incurred claims to choose/calibrate the severity distribution
Bypass the need to calculate IBNER by using only closed claims, with some correction mechanism to eliminate the bias that derives from assuming that claims will settle regardless of size
Murphy-McLennan method: use historical development of claims and a chain-ladder-like methodology to identify patterns in the development of claims
Generalised linear modelling to identify the factors that drive the IBNER development factors
Other sensible methods that have been proposed or are used

(ii)

<p><i>Ignore IBNER and use all incurred claims</i></p> <p>(+) The method is very simple, and only requires a snapshot of the loss run at a given point in time – no need for the full historical development of claims</p> <p>(+) It will work fine if there is no systematic over- and under-reserving or if only a few claims are open</p> <p>(+) It is the only option when there are not enough loss data to produce a credible statistical analysis of IBNER, and when no industry IBNER factors are not available.</p> <p>(-) If there is systematic over- or under-reserving, the results will inevitably be biased</p>
<p><i>Use closed claims only</i></p> <p>(+) Does not use loss estimates and therefore the input to the analysis is more certain</p> <p>(-) It doesn't use all data and therefore "conceals" evidence</p> <p>(-) It cannot be used without a bias-correction mechanism to cater for the fact that large claims are the ones that are more likely to take longer to settle</p>
<p><i>Murphy-McLennan (chain-ladder on historical reserves)</i></p> <p>(+) Allows to use all data, even the estimated amount</p> <p>(+) It is relatively simple, as development factors can be calculated, e.g., by chain ladder techniques</p> <p>(-) It requires more data than the two methods above</p> <p>(-) In its simplest implementation, it does not take into account the fact that losses of different size and different outstanding percentage may have different development, and therefore the estimate may be too rough.</p>
<p><i>Identify trends and adjust for IBNER by using GLM-like techniques</i></p> <p>(+) Allows to use all data, even the estimated amount</p> <p>(+) Allows to consider the impact of different factors such as loss size and outstanding percentage on the way the loss will develop</p> <p>(-) It is a relatively complex method</p> <p>(-) The method only work with a significant amount of data</p>

(iii)

First of all, exclude all years for which the O/S ratio is = 0%....				
... since these claims cannot develop any further except in case of re-opening				
These are the remaining years:				
d=2	d=3	Factor 2->3	O/S ratio	
1,041,681	1,332,408	1.28	99%	
48,001	48,001	1.00	100%	
30,494	2,777	0.09	98%	
149,398	228,689	1.53	42%	
80,387	80,387	1.00	65%	
280	6,819	24.36	81%	
1,355	1,355	1.00	58%	
The straight average of the IBNER incremental factor is $(1.28+...+1)/7=4.32$				
The weighted average is $(1,332,408+...+1,355)/(1,041,681+...+1,355)=1.26$				
The median is the 4 th element in set of the sorted values: (0.09, 1, 1, 1, 1.28, 1.53, 24.36)				

(iv)

<i>Straight average</i>
(+) It is straightforward and takes full account of all the changes from one period to the other (unlike the median that is only affected by the values in the middle of the distribution)
(-) It is very sensitive to outliers, e.g. the 24.36 value
(-) It gives the same weight to changes that apply to hugely different money amounts: e.g. the change from £280 to £6,819 is minor in monetary terms but huge in relative terms (x 24.36), and it increases the average up to 4.32 with the result that a large factor may be applied to much larger amounts.
<i>Weighted average</i>
(+) Straightforward and being money-weighted it is less sensitive to shaky menial amounts
(-) It's still sensitive to individual outliers
<i>Median</i>
(+) Very robust statistic, insensitive to outliers
(-) Since most open claims in many circumstances tend to remain unchanged from one period to the other, the median will often be 1 and there will be no IBNER even in situations where, on average, there is a systemic impact.

Question 2

The IBNER-adjusted loss amount is given by Adjusted(ID) = Loss@current_position(ID)*Factor(d-->onwards) where d = Year of current position - Year of reporting + 1 (alternatively, a clear example of detailed calculation)							
Actual calculations:							
Loss ID	Loss amount @ current position	Outstanding percentage	Year of loss	Year of reporting	Year of current position	Relevant factor	Adjusted loss amount
1	10,000	50%	2009	2012	2013	1.32	13,167
2	30,000	15%	2010	2010	2013	0.95	28,500
3	25,000	0%	2010	2011	2013	1.00	25,000
4	45,000	10%	2008	2009	2013	1.00	45,000
5	7,000	100%	2011	2013	2013	2.44	17,080

The losses with outstanding percentage = 0% should *not* be adjusted

Chapter 16

Question 1

Advantages of a stochastic model

It allows to give a full distribution of results with all the percentiles needed rather than a rough idea of the mean and the standard deviation (some percentiles may also be calculated with burning cost analysis, but this is limited by the number of years)

It deals with IBNR properly – burning cost results can be modified to include an allowance for IBNR but this requires information external to the client

It is indeed a *model* and can therefore make predictions outside of the scope of the actual losses experienced.

It should give a better idea of the volatility of the results.

Disadvantages of a stochastic model

More difficult to communicate to non-actuarial audiences – underwriters/clients will often be more familiar with the workings of a BC model, and more easily convinced by its results.

It is more complicated and takes longer to develop

It needs more data than a BC model, in which no fitting is involved

One can easily have an exaggerated confidence in the results of a stochastic model and a hazy idea of its uncertainties, therefore thinking that it is possible to say something meaningful about, e.g., very high percentiles that go far above the limits of the actual experience (this is the flipside of the second advantage!).

Unless we are very careful about it, there is always the risk of overfitting, e.g. by selecting models that replicate very well the observed results but with little predictive value

Chapter 17

Question 4

(i)

An AAD is the maximum amount which is retained by the insured in policies that have an each-and-every loss (ranking) deductible.
Its main purpose is providing additional certainty as to the overall amount that the insured will have to retain in a bad year...
... although this will not protect the insured from very large losses that exceed the policy limit (£25m in this case)

(ii)

<p>The AAD works as follows: if X_1, X_2, \dots, X_n are the losses incurred in a given year, the amount retained <u>below the deductible</u> becomes:</p> $\text{Total retained below the deductible} = \min(\text{AAD}, \sum \min(X_j, \text{EEL}))$ <p>[alternatively, explain the procedure in words, either in isolation or while doing the calculations below]</p> <p>E.g. if EEL = £100k, AAD=£300k, and we have four losses $X_1=£80k, X_2=£120k, X_3=£70k, X_4=£150k$, then Total retained below EEL = $= \min(\text{£}300k, \min(\text{£}80k, \text{£}100k) + \min(\text{£}120k, \text{£}100k) + \min(\text{£}70k, \text{£}100k) + \min(\text{£}150k, \text{£}100k)) =$ $= \min(\text{£}300k, \text{£}80k + \text{£}100k + \text{£}70k + \text{£}100k) = \text{£}300k$</p>
--

(iii)

There is more than one acceptable solution to this. However, the standard solution involves the use of MC simulation, and works as follows:
<ol style="list-style-type: none"> Run N simulations combining the frequency and severity model. Each simulation j will have a number of losses $X(1), \dots, X(n(j))$ Produce the uncapped retained amount for each simulation by capping every loss to the deductible: $X(1, \text{ret}) = \min(X(1), \text{EEL}), \dots, X(n(j), \text{ret}) = \min(X(n(j)), \text{EEL}) \rightarrow$ $\text{Retained_Uncapped}(j) = X(1, \text{ret}) + X(2, \text{ret}) + \dots + X(n(j), \text{ret})$ Sort the results ($\text{Retained_Uncapped}(1), \dots, \text{Retained_Uncapped}(N)$) in ascending order, yielding the empirical uncapped retained loss distribution Pick the 95th percentile of the empirical retained loss distribution, c^* AAD must be $\geq c^*$ in order to abide by the insurer's constraints

(iv)

One can use the results of the simulations in Step (iii) to calculate this effect, and changing Step b as follows Produce the uncapped retained amount for each simulation by capping every loss to the deductible, and then imposing a cap equal to AAD:
--

$X(1,ret)=\min(X(1),EEL),\dots X(n(j),ret)=\min(X(n(j)),EEL) \rightarrow \text{Retained_Capped}(j) = \max(AAD,X(1,ret)+X(2,ret)+\dots+X(n(j),ret))$
Calculate the average across all simulation of Retained_Capped , Avg_Retained_Capped
The difference between $Z = \text{Retained_Uncapped} - \text{Retained_Capped}$ gives the increase in the expected ceded losses to the insurer

(iv)

<p>One can use the results of the simulations in Step (iii) to calculate this effect, and changing Step b as follows</p> <p>Produce the uncapped retained amount for each simulation by capping every loss to the deductible, and then imposing a cap equal to AAD:</p> $X(1,ret)=\min(X(1),EEL),\dots X(n(j),ret)=\min(X(n(j)),EEL) \rightarrow \text{Retained_Capped}(j) = \max(AAD,X(1,ret)+X(2,ret)+\dots+X(n(j),ret))$
Calculate the average across all simulation of Retained_Capped , Avg_Retained_Capped
The difference between $Z = \text{Retained_Uncapped} - \text{Retained_Capped}$ gives the increase in the expected ceded losses to the insurer

(v)

One thing you can do is to calculate the retained loss amount (uncapped) for all years...
... making all the necessary adjustments, e.g. for claims inflation, IBNR, etc.
... then calculate the "burning cost" for each year, i.e. retained loss amount divided by exposure, thus obtaining ≤ 10 values (the "<" is in case some of the years need to be excluded)
Fit a distribution (e.g. a normal, or lognormal) distribution to these different values. This distribution models the year-on-year volatility of the burning cost
Pick the 95 th percentile of this distribution, c^*
The aggregate AAD must be $\geq c^*$

(vi)

The burning cost approach is much simpler and less time consuming than the stochastic loss model approach
It is difficult to justify the use of any particular statistical distribution (e.g. normal) to model the year-on-year volatility of the burning cost , based on only ≤ 10 points (model uncertainty)
A single year which is out of line may hugely increase the year-on-year volatility and produce estimates of the 95 th percentile which are of a different order of magnitude with respect to the correct one. As a result, the estimate based on the burning cost may simply be unreliable (parameter uncertainty)
The burning cost approach does not take into account separate frequency/severity trends...
... and is normally purely retrospective...
... and hugely affected by the absence/presence of large losses

Question 5

(i)

We assume that the new model sold at Christmas will have a similar failure rate to that of the model for which a study has recently been undertaken
We also assume that the proportion of cameras sold in different countries of Europe is the same as in the study (otherwise, if more cameras are sold in countries with a higher failure rate, our 5% estimate would not be valid even if the model hadn't changed)
The expected frequency is $E(N)=25,000 \times 5\% = 1,250$ The expected severity is $E(X)=€25$ The expected losses are $E(S)=E(N) \times E(X) = €31,250$

(ii)

The most straightforward way is to notice that the number of failures follows (in case the failures are independent) a binomial distribution with $N=25,000$, $p=5\%$...
therefore the probability of having k failures is $\Pr(k) = \binom{N}{k} p^k (1-p)^{N-k}$
And the total cost of having k failures is $C_{tot} = k \times C_U$ where $C_U = €25$ is the unit cost of replacement.
The 90 th percentile corresponds to a number of failures K such that $K^* = \max(K \sum_{k=1}^K \binom{N}{k} p^k (1-p)^{N-k} \leq 0.9)$
And the total cost corresponding to the 90 th percentile is therefore $C_{tot} = K^* \times C_U$
Alternatively, one can use a Monte Carlo simulation as follows: (i) "Scan" all 25,000 cameras (ii) For each camera, generate a pseudo-random number and if this number is $< 5\%$ (the probability of failure), consider that this camera has a failure (iii) Count the number of failures across all cameras (iv) Repeat the scan a large number of times, e.g. for 10,000 simulations, and for each calculate the number of failures (v) Sort the number of failures for each simulation in ascending order, and pick the 90 th percentile, which corresponds roughly to the 9,000 th value in ascending order.

(iii)

The failure rate does not refer to the model which will be sold for Christmas, but to a similar model. (How similar that is we don't know, but different models of a device will have different failure rates, and some of them could also have systemic ("epidemic") problems.
We are given no reassurance as to whether the proportion sold across the various territories for MB4 will be the same as for MB2 and this might have further impact on the failure rate assumption.

Chapter 18

Question 1

(a)

Parameter uncertainty –

(the parameters of your Negative Binomial and Lognormal will be calculated perhaps with MLE or similar methods and a standard error is associated with these estimates)

Your estimate of parameter uncertainty depends on the estimation method.

As an example, in MLE you can use the inverse of Fisher’s matrix

Bootstrap is also a good method which can be used with different estimation methods...

...It is possible to re-run your 99.5% estimate using different parameter estimates and get an idea of the volatility.

Data uncertainty –

300 out of 1,000 losses are only estimates. Depending on how they are actually settled, the parameters of your model will need to be changed

Claims inflation is also only an estimate and depending what value we use, all losses will change and the parameters will change again.

One way of estimating the effect of data uncertainty is that of assigning a prior distribution to your uncertain data and to inflation and re-run your model many times assessing the volatility of the 99.5th percentile (parametric bootstrap)

Model uncertainty.

This is a tricky one. The point is that your Negative binomial/Lognormal model may not be the right one even if it’s the best one you’ve been able to find!

It’s difficult to measure model uncertainty unless you have a prior probability of models being right.

What you can do is assess the goodness of fit via cross-validation or on a hold-out sample to have a fair idea of how well your model is performing.

You might also want to compare the empirical model (drawing samples from the actual losses) to the NB/Lognormal model to get an idea of the order of magnitude of the difference in the 99.5th percentile between one model and the other.

Or you can compare the results for different models...

Simulation error –

Of course even when you’re confident of a model, of your data, of your parameters, the 99.5th percentile will vary from one simulation to the other.

Run your simulation many times to ascertain the volatility in the results.

Process uncertainty – this is not an issue here.

(b)

The risk manager’s concerns on your 1 in 200 estimate are more than justified!

It is difficult to make accurate predictions in the tail of the distribution...

...especially using a lognormal distribution whose parameters will be mostly driven by small to medium losses...

... whereas it’s the large losses which most likely will affect the behaviour at high percentiles.

(c)

You can at least model the tail separately using extreme value theory (i.e. model the tail as a GPD above a certain threshold).

It might be helpful to use market data (e.g. data from other companies, or reinsurance data) to have a more realistic idea of how the losses scale up in the large losses region.

You could then use the client's data to model the body of the distribution, and the market data to model the tail.

Another option is to use credibility weighting, if you have some benchmark

Chapter 19

Question 2

(i)

The formula for the technical premium, once the net reinsurance premium and the commission is stripped out, becomes

$$\text{Technical premium} = \frac{\text{Expected losses}}{(1 - \text{Expenses}\%) \times (1 - \text{Profit}\%) \times (1 + r)^t}$$

(reasonable alternatives will also be accepted).

Where r is the risk-free rate for the relevant weighted payment time

And t is the weighted payment time

The risk-free rate can in our case be approximated as the average of the yields for a 5-year gilt and a 7-year gilt: $\text{Yield}(6y) \sim (\text{Yield}(5) + \text{Yield}(7)) / 2 = 1.17\%$

However, note that technically a “risk-free” rate is usually associated with triple-A government bonds, and UK gilts have *[at the time of writing this]* recently lost that status. However, this remains a good approximation.

The technical premium becomes in our case:

$$\text{Technical premium} = \frac{\text{£1.2m}}{(1 - 28\%) \times (1 - 10\%) \times 1.0117^6} = \text{£1.73m}$$

(ii)

The net cost of reinsurance can be written as:

$$\text{Net cost of reinsurance} = \text{Reins. premium} - \text{Expected recoveries}$$

Based on the information given in the text,

$$\text{Net cost of r/I} = \text{£800k} - \text{£550k} = \text{£250k}$$

The revised technical premium is therefore

$$\text{Technical premium} = \frac{\text{£1.2m} + \text{£0.25m}}{(1 - 28\%) \times (1 - 10\%) \times 1.0117^6} = \text{£2.35m}$$

(iii)

The technical premium is not the premium actually charged, but the premium calculated according to the actuarial analysis and the company guidelines on expenses, profit, etc.

The actual premium charged is the result of a strategic decision...

... which involves considerations of the premiums charged by competitors, and the desire to maintain market share...

... the relationship with other products (e.g. aviation liability is likely to be sold as a package with aviation hull, and it is the combined performance of the product which matters)...

... regulations and other external elements (this is unlikely to be a strong factor for aviation liability)

Confidence in the company's ability to achieve a much better investment return than the risk-free rate

Since aviation liability has losses with very large impact but that happen with very low frequency, and depending on the incentives system for underwriters, the underwriter may be focusing on short-term premiums, hoping that no losses will materialise for the next few years...

... and the availability of cheap reinsurance may limit the risks of such a strategy

Pricing optimisation techniques, which have been developed in a personal lines context, are not applicable to a product like aviation liability.

Chapter 20

Question 3

(a)

A losses occurring policy is one where the reinsurer assumes responsibility for the claims which hit the reinsurance layer (subject to limits, deductible, reinstatements, etc) and that have occurred to any of the original policies of the insurer during the (reinsurance) policy period, irrespective of when the policies incepted.

(b)

A reinstatement is the restoration of full cover following a claim.

Reinstatements are used by reinsurers to limit the amount they are liable to pay to the reinsured, and therefore achieve more certainty in their results.

At the same time, they may benefit the *reinsured* in that they help bring down the price of the reinsurance contract.

(c)

Since the size of the layer is £3m, and you can reinstate it twice, that means that overall you have an aggregate limit of three times the size of the layer, for a total of £9m.

(d)

The two losses impact on the £3m xs £5m layer by £2.5m and £3m respectively.

The sum is £2.5m+£3m = £5.5m, which is less than three times the layer size and is therefore fully recovered.

The first £3m of losses (£2.5m of the first one and £0.5m of the second loss) trigger a full reinstatement at zero premium. The rest (£2.5m) triggers a reinstatement premium equal to £2.5m/£3m x 50% of £400k =£167k.

Therefore overall the premium paid is £400k (upfront premium) + £167k (a portion of the second reinstatement @ 50%) = £567k

(e)

The purpose of the index clause is to counter the effect of claims inflation on claims that take long to settle, so to decrease the impact of such a loss on the reinsurer and decrease the reinsurance premium charged to the insurer.

Since this is a property risk XL policy, most losses are settled relatively quickly and the effect of claims inflation is much smaller than for liability policies, for which the index clause was conceived. Therefore normally you only use the index clause in liability XL policies.

Question 4

(a)

- The payment/settlement pattern needs to be included...
- ... since it is needed to adequately assess the effect, e.g., of the indexation clause

(b)

All historical claims over the last ten years will need to be revalued to the mid-term of the policy, which is 1/7/2013. A claim occurred on 1/1/2003 (the earliest possible date) would need to be revalued T=10.5 years.

The reporting threshold R must therefore be such that

$$R \leq \frac{AP}{(1+r)^T} = \frac{£1m}{1.1^{10.5}} = £368k$$

An acceptable threshold is any round number below that level without being unjustifiably low: £300k or £350k are reasonable values.

If R were higher than £368k, say £400k, then, e.g., a claim of £380k occurred on 2/1/2013 wouldn't be included in the data set, despite the fact that its revalued amount would be higher than £1m and would therefore be relevant to the treaty.

This would create a bias by which lower claims (those just above £1m) are under-represented in the loss severity distribution.

Note that the layer limit (£4m) is not relevant to the calculations above.

Question 5

Assume that

- The premium is earned uniformly over the year
- All premium is annual
- The premium in 2005 is the same as the premium in 2006, or some other suitable assumption
- There is no IBNER
- There is no change in the risk profile
- There are no changes in the cover (policy terms and conditions)
- The rate change is applied at the beginning of the year only
- No correction necessary for large losses
- Claims inflation is 5% over the period (already in the text)

- No rate change in 2005 (already in the text)
- Same proportion of EL and PL (already in the text)

There are several modifications to the claims incurred:

- Claims must be revalued by 5% per annum, starting from 2012 (100% - no revaluation) → Column “Revaluation factor”
- Claims must be adjusted by IBNR, using the IBNR factors provided in the text

As for the premium, this must be changed to account for rate changes. Rate changes are rating actions that the underwriter takes *across the board* (i.e., on average for all clients) to compensate for the increased/decreased risk from year to year (e.g. because of claims inflation), or for changed market conditions.

- The written premium must be adjusted to take into account the cumulative rate change, in a multiplicative fashion: e.g. if from 2008 to 2011 there were three rate changes: 7%, -2%, 3%, the cumulative effect is 107% x 98% x 103% - 1 = 8.006% (you don’t need to be that accurate – it’s just to show the principle behind the calculation)
- The earned premium can be calculated (given the above assumption on the uniformity of the way in which premium is earned over the year) as the average of the premium in the present year and the previous year: $EP(n) = (WP(n-1) + WP(n))/2$
- It is important to apply the rate changes to the written premium and not to the earned premium!

The loss ratio can then be calculated as

$$LR = (\text{Total claims incurred over 2006-2011}) / (\text{Total premium earned over 2006-2011}) = 91.5\%$$

The calculations are shown in the table below.

Policy year (incepting @1/1)	Premium written during the year (all annual policies)	Rate change w/ respect to the previous year	Claims incurred during policy year	Develop ed	Revaluati on factor	Revalued claims	Projecte d claims	Cumul rate change - factor	Premium written - adjusted	Premium earned
2006	1.45	N/A	1.53	100%	134%	2.05	2.050	116.808	1.69	1.69
2007	1.5	3%	1.15	100%	128%	1.47	1.468	113.406	1.70	1.70
2008	1.8	5%	1.25	100%	122%	1.52	1.519	108.006	1.94	1.82
2009	1.75	7%	1.67	95%	116%	1.93	2.035	100.940	1.77	1.86
2010	1.9	-2%	1.05	80%	110%	1.16	1.447	103.000	1.96	1.86
2011	1.95	3%	0.41	30%	105%	0.43	1.435	100.000	1.95	1.95
2012 (estim)	2	0%		0%	100%	-		100	2.00	1.98

Summary stats as at 31/12/2011

LR = 91.5%

Chapter 21

Question 7

(i)

- d represents the ratio between the deductible imposed and the sum insured or maximum possible loss
- $G(d)$ represents the percentage of risk retained by the reinsured if a deductible d is imposed

(ii)

The larger the concavity, the larger the impact of a deductible.

The closer the curve is to the diagonal, the closer we are to a situation where each loss is a total loss

(iii)

The expected losses to a layer $(D, D+L)$ for a property with sum insured/MPL equal to M are given by

$$E(S^{D,L}) = (G(d+l) - G(d)) \times E(LR) \times P$$

Where $d=D/M, l=L/M$

The calculation for the first group of properties (those with sum insured between £4.5m and £5m) is shown below. The calcs for the other groups are similar.

- Assume $M=£4.75m$ (mid-point) for all properties
- $d = D/M = £2m/£4.75m = 42.1\%$
- $d+l = (D+L)/M = £5m/£4.75m > 1 \rightarrow$ replace with 1
- $G(d) = G(42.1\%) \sim 87\%$ (from the chart)
- $G(d+l) = G(100\%) = 100\%$
- Expected losses (1) = $(G(d+l) - G(d)) * LR * P \sim 13\% * 66\% * £1.6m = £138k$

The total expected losses are simply given by the sum of the three bands:

$$\text{Tot exp losses} = EL(1) + EL(2) + EL(3) = £138k + £169k + £67k = £374k$$

Chapter 22

Question 3

(a)

ILFs (increased limit factors) are multiplicative factors that give the ratio of the premium at a higher limit L based on the knowledge a premium at a basic limit B ($B < L$).

ILFs can be derived given the underlying loss severity distribution $F_X(x)$ based on the following formula:

$$ILF(x) = LEV(x)/LEV(b)$$

Where $LEV(x)$ is the limited value of the severity distribution ($LEV(x) = E(\min(x,X))$)

(b)

ILFs are used typically to price liability business...

... to understand the effect of raising the policy limit, often much beyond the normal loss experience available...

They are often used in reinsurance...

... especially to price excess-of-loss business

(c)

Exposure curves are based on relative losses (i.e. as a proportion of a total loss), whereas ILFs are based on absolute amounts.

Exposure curves are for risks that have an upper bound, whereas ILFs are for risks for which there is no theoretical limit.

(d)

Since ILFs are based on the LEV for the underlying loss severity distribution, an ILF can be derived by building an empirical distribution based on a database of past losses...

... and applying the relationship $ILF(x) = LEV(x)/LEV(b)$ repeatedly

Some data-related issue in the construction of an ILF curve are as follows:

- Some large losses may be truncated as a result of policy limits
- Data need to be trended appropriately, as in experience rating (not a major problem, but it involves an additional uncertainty)
- Many claims will not be closed, and one needs a strategy on how to deal with those, as they can't simply be left out without further adjustments, and there are dangers in using the reserve estimate rather than the closed amount

(e)

We assume that the effect of inflation is uniform over the whole range of losses, and that therefore the shape of the ILF curve doesn't change because of inflation.

However, this may not always be the case: e.g., large losses often have higher inflation than smaller ones

$$ILF(\pounds 500k) = \frac{ILF\left(\frac{\pounds 500k}{1.05}\right)}{ILF\left(\frac{\pounds 100k}{1.05}\right)} = \frac{ILF(\pounds 476,190)}{ILF(\pounds 95,328)} = \frac{1.773}{0.953} = 1.860$$

$$ILF(\pounds 1m) = \frac{ILF\left(\frac{\pounds 1m}{1.05}\right)}{ILF\left(\frac{\pounds 100k}{1.05}\right)} = \frac{ILF(\pounds 952,381)}{ILF(\pounds 95,328)} = \frac{2.281}{0.953} = 2.393$$

(f)

Old year: $P = (2.285 - 1.785) E = 0.500 E$

New year: $P = (2.281 - 1.773) \times 1.05 E = 0.533 E$

Change: $0.5333 E / 0.500 E = 1.067 (+6.7\%)$

Chapter 23

Question 1

(i)

Occurrence basis

All losses occurring during the policy period will be recovered from the insurer, regardless of when they are reported

Claims made

All losses reported during the policy period will be recovered, regardless of occurrence date, as long as they are after the retroactive date.

(ii)

Since the policy is claims-made, the relevant exposure associated with the claims reported in one year y is not the turnover in that same year only

but also the turnover in all years $y-1, y-2, \dots$ from which a claim may ultimately originate

i.e., all years in which the legal advice has been given...

... and that still have an impact on the claims experience y , given the reporting pattern

In our case, the reporting pattern is 25% (y), 35% ($y+1$), 40% ($y+2$), and therefore the adjusted exposure in year y is given by

$$\text{Adj_Exp}(y) = 0.25 \times \text{Exp}(y) + 0.35 \times \text{Exp}(y-1) + 0.4 \times \text{Exp}(y-2)$$

We should also check that we don't count any exposure for years before the retroactive date

However, with our assumptions the retroactive date does not affect our calculations, because the earliest year for which we have recorded claims experience is 2008 which is affected only by 2008, 2007, 2006

(iii)

The adjusted exposure for every year can be found in the following table:

<i>Year</i>	<i>Exposure (turnover, revalued), \$m</i>	<i>Claims reported (already revalued), \$m</i>	<i>Exposure (adjusted)</i>
<2008	110		
2008	110	1.80	110
2009	100	1.72	107.5
2010	80	1.66	99
2011	50	1.07	80.5
2012	50	1.13	62
2013 (est)	30		45

The burning cost estimate for 2013 is given by:

$$\begin{aligned}
 BC &= (\text{Claims}(2008)+\dots+\text{Claims}(2012))/(\text{Adj_Exp}(2008)+\dots+\text{Adj_Exp}(2012))*\text{Adj_Exp}(2013)= \\
 &= (1.80+1.72+1.66+1.07+1.13)/(110+100+80+50+50)*45 = \\
 &= \$0.724\text{m}
 \end{aligned}$$

(iv)

Extra assumptions

- Turnover period and policy period match
- All claims reported have already been quantified, so that there are no RBNQ (reported but not yet quantified losses)
- There'll be no further adverse or favourable development of the claims already reported

(v)

The main problem is that all losses for advice given prior to 2013 and not reported by the end of policy year 2012 is neither going to be covered by the previous claims-made arrangement nor by the new occurrence-basis arrangement!

If X still wants to go ahead with the change, it will need to buy tail cover...

... which covers all losses occurred *prior* to 2013 but reported *in 2013 and beyond*

Question 2

(i), (ii) See Section 23.5.3.5

(iii)

Claimant's point of view

The lump sum approach achieves finality and provides the claimant with flexibility both in terms of how he invests the lump sum and how and when he applies it for his needs.

However, with lump sums there's a risk that the claimant may exhaust the lump sum through overspending, poor investment (investment risk), because of excessive inflation (inflation risk) or simply living longer than anticipated (longevity risk), he can be forced to fall back on the resources of the state.

In addition, lump sum awards provide no opportunity for adjustment should a claimant's needs change save for on the rare occasions when damages are awarded on a provisional basis

The advantages/disadvantages of PPOs are of course the flipside of the advantages/disadvantages for lump sums above.

Insurer's point of view

PPOs lead to uncertainty in final payment because of the uncertainty of future life expectancy (longevity risk) and the possible changes in the claimant's needs which may lead to a revision of the periodic payment

Also, the insurer assumes the investment risk and inflation risk that in the case of the lump sum are with the claimant.

PPO lead to an increase in costs (more capital to set aside, more claims handling costs, higher costs of reinsurance, need to purchase annuities...)

One would be hard-pressed to find an advantage of PPOs for insurers! However, it may turn out that the overall payments paid periodically is lower than the estimated initial lump sum

Question 3

You are a quantitative analyst for Company X.

Company X sells cheap digital cameras which are sold at €50 in the European market and have a production cost of €20. The cameras are offered with a one-year warranty that offers a free replacement for failures during the warranty period.

A new model (MB4) of camera with the same price as the rest of the range is going to be sold in the 2013 Christmas period in and 25,000 sales across Europe are expected.

A recent study of the probability of failure suggests that the failure rate for a similar camera model (MB2) was 5% during the first year, although some differences among different European territories were noticed.

You can assume that:

- You can only get one replacement a year.
- The overall replacement cost is €25 (=production cost + expenses including mailing expenses).
- If a camera is replaced, the warranty expiry date remains the original one: it does not start again from the date when the camera is received.

P. Parodi – Pricing in general insurance – Solutions to selected problems

- (i) Calculate the expected cost of providing warranty cover for the lot of MB4 models sold during Christmas. State clearly all the assumptions you make.
- (ii) Describe how you would calculate the 90th percentile of the statistical distribution of the costs (you don't need to carry out the actual calculations, but to only outline the methodology clearly enough for one to follow). State clearly all the assumptions you make.
- (iii) What are the *main* uncertainties around the expected cost calculated in (i)?

Question 4

(i)

The following items are typically covered by a PA policy:

- Loss of limb
- Loss of eyesight (one eye)
- Loss of eyesight (both eyes)
- Death
- Temporary disablement

Example of schedule

- Loss of limb, £20,000
- Loss of eyesight (one eye), £10,000
- Loss of eyesight (both eyes), £20,000
- Death, £25,000
- Temporary disablement, £250/wk

The exact figures included are not important; what is being tested here is the understanding that this is a fixed benefits policy with a fixed amount for each type of loss.

Also, there should be some consistency with what said in the coverage.

(ii)

Two obvious examples are:

Car insurance
Travel insurance

The personal accident element in car insurance is there to provide compensation for specific types of bodily injury suffered by the driver as a consequence of a motor accident.

The personal accident element in travel insurance is there to provide compensation for specific types of bodily injury suffered by the policyholder while travelling

(iii)

Most non-life insurance works on the principle of indemnity...

... by which the policyholder is restored to the financial situation (s)he was before the loss

However, it is not possible to restore someone's position to that prior to the loss for bodily injury or death (although courts may quantify things such as loss of earning or pain and suffering in a liability context), therefore a fixed benefit is decided beforehand for any type of loss covered, much in the same way as this is done, e.g., in life insurance

(iv)

We need to assume constant exposure, i.e. constant number of policies written between 2008 and 2013

... and an unchanged risk profile

... and that all deaths have been reported up to year 2012

... and there's no trend in the number of deaths

The total loss expected for policies incepting in 2013 will be driven by two elements:

<ul style="list-style-type: none"> - The expected number of claims - The severity of each claim
The expected number of claims can be estimated e.g. by calculating the average number of claims over 2008-2012, which is 2,463
The expected severity of the claims is not an issue – it is set in the contract at £1.6m
The expected losses is therefore given by $2,463 \times £1.60m = £3.94m$

(v)

The most common frequency models used for frequencies are binomial, Poisson, negative binomial
The variance/mean ratio can help discriminate, because for the bin distr it is <1 , for the Poisson distr is $=1$, for the negative bin is >1
In this case, the variance to mean ratio is 8.05
Hence, the negative binomial distribution should be chosen

Chapter 24

Question 1

(a)

Catastrophes are by definition very rare events, and therefore the burning cost associated with them is likely to be zero over 10 years of experience. Analogously for a frequency/severity model, the loss experience will normally be nil or very small over that period.

If for some reason the burning cost/the loss experience is significant it means that a catastrophe has happened and the burning cost may exaggerate the true long-term cost of risk

More in general, experience-based analysis works best for low severity, high frequency events, and estimating the impact of events that might have a return period of, say, 1 in 250 years, based on the loss experience over a period of 10 years results in a large uncertainty.

(b)

A catastrophe model deals with the scarcity of data by referring to a larger, external data set of known historical events...

... and by using our scientific understanding of the underlying perils to create possible future catastrophic events...

... including events that have never been observed historically.

The model also calculates the impact of these events on the insured portfolio, ultimately creating a list of realistic losses with associated probabilities that are used instead of the actual experience to estimate the cost of insurance.

(c)

First of all, note that different publications will have different definitions of the modules and the number of modules change as well, e.g. if we consider the exposure data a module in itself or one of the inputs to the modules. In the following, we use the definitions of the lecture notes.

Hazard module

This includes two steps: (i) the generation of events and (ii) the assessment of the level of physical hazard for each generated event.

i. Event generation set

In this step, a catalogue of potential catastrophe events is generated, based on a collection of historical event data for the variables above and an analysis of the probability distributions for these variables.

ii. Hazard assessment

The hazard component of catastrophe models assesses the level of physical hazard across a geographical area at risk. E.g. for hurricanes, a model calculates the strength of the winds around a storm, considering the region's terrain and built environment.

Vulnerability module

Vulnerability can be defined as the degree of loss to a particular system or structure resulting from exposure to a given hazard (often expressed as a percentage of sum insured).

The vulnerability component calculates the amount of expected damage to the properties at risk. Vulnerability functions are region-specific, and vary by a property's susceptibility to damage from earthquake ground shaking or hurricane winds.

Parameters defining this susceptibility include a building's construction material, its occupancy type, its year of construction, and its height. In catastrophe models for insurance applications, different

P. Parodi – Pricing in general insurance – Solutions to selected problems
vulnerability curves are used to estimate damage for a structure, its contents, and time element coverage such as business interruption loss or relocation expenses.

This module also includes critical estimates of uncertainty around expected damage (i.e., standard deviations). Together, the stochastic event, hazard and vulnerability modules comprise what is traditionally known as a probabilistic risk analysis.

In state of the art models, the vulnerability module often includes also an '**inventory**'. The inventory is a database that contains estimates of all building attribute mixes possible in the form of percentage weights for a given location and occupancy (ideal for scarce client data).

Financial Analysis module

This module translates physical damage into total monetary loss. Estimates of retained and ceded losses are then computed by applying policy conditions (such as deductibles and limits) to the gross loss estimates.

The outputs of the financial module are the Event Loss Tables (ELT), which give the detail of the event-by-event losses. There is a different ELT for each insurance structure considered, and a *gross* ELT (no insurance structure).

(d)

Event	Rate	Expected Loss	Standard Deviation	Exposure
1 – Northern San Andreas 6.5	0.01	1,500,000	800,000	5,500,000
2 – Calaveras 6.5	0.01	3,000,000	2,000,000	15,000,000
3 – Hayward 7.0	0.02	6,500,000	5,000,000	50,000,000
4 – Northridge 6.5	0.03	8,000,000	6,000,000	90,000,000
5 – Southern San Andreas 5.0	0.03	10,000,000	7,000,000	95,000,000

The first column gives the event description (in this case, it's earthquakes with their Richter scale factor).

The second column gives the probability of that event happening.

The third column gives the expected value of the gross damage of that event.

The fourth column gives the volatility around that value, which is also called secondary uncertainty

The fifth column gives the total amount exposed to that event – i.e. it is an upper bound on the possible loss value.

(e)

The overall yearly rate is simply given by the sum of all the rates in the second column of the ELT: in the case shown above, this is $r = 0.01 + 0.01 + 0.02 + 0.03 + 0.03$

Losses can be simulated by sampling events from the ELT with probability proportional to the rate (second column), and with severity drawn from a distribution with mean = to the mean in the third column, and standard deviation = to the SD of the fourth column

This is usually achieved by using a Beta distribution with values between 0 and the total exposure value (fifth column), mean = column 3, std dev = column 4

(f)

First and foremost, terrorist attacks are organised and carried out by people and therefore are inherently less predictable (almost by design), in the same way that the stock exchange is less predictable than the weather.

Whereas the weather has no particular interest in being unpredictable, and does not try to game the systems, humans naturally do.

It is very difficult for models to take these aspects into account.

Information on terrorism becomes quickly obsolete, as different issues which cause friction between communities resolve or are exacerbated over the years, and completely new issues/terrorist groups arise.

Whereas, say, California is bound to remain an earthquake-prone area for the foreseeable future, a region in the world may become all of a sudden more or less risk depending on specific events.

Underwriting terrorism in a particular region requires up-to-date knowledge of the political issues concerning that region – again, this is difficult to model and update continually.

The severity distribution of terrorist events may change quite rapidly as new warfare technology (eg bacteriological weapons) on the one hand and new risk control mechanisms (eg airport security) are introduced. In the case of nat cats, changes are more likely to be piecemeal, eg in response to climate change.

Terrorist risk may under certain conditions be systemic and geographically widespread. E.g., events in the Middle East may trigger threats to Isreal interests all over the world. This is partly the case

As in the case of nat cats, there may be seasonalities for terrorist risk, but of a more complicated kind, e.g. corresponding to anniversaries.

The number of issues involved in building a terrorist model is larger and more varied than in the case of nat cats: it involves the knowledge of many different cultures and political situations, and basically a constant attention to news events. For this reason, it will be very expensive to maintain.

Chapter 25

Question 3

Formula:

$$Z = \frac{\sigma_h^2}{\sigma_h^2 + \sigma_c^2}$$

-
- σ_c is the standard deviation of the client risk premium estimator and measures the (lack of) credibility of the client's experience, or in other words the uncertainty by which the risk premium of the client is calculated
-
- σ_h is the standard deviation of the risk premiums of different clients in the market – in other terms, it represents the heterogeneity of the market risk and therefore its (lack of) relevance for the client we're analysing
-

Properties:

- Z is between 0 and 1
-
- which is sensible as the premium charged will be somewhere between the premium calculated from the particular risk and the premium charged when looking at all risks
-
- Z is a decreasing function of σ_c
-
- The more uncertain the risk premium calculation based on the client's data only is, the less credible the result obtained from looking at just that particular client and so the bigger the need to look at external data
-
- Z is an increasing function of σ_h
-
- The more variation between risks the less inclined we are to use data from other risks as risks are not homogeneous
-

Question 4

(a)

Experience rating					
- List of all claims over the last N years (N=10 or so If possible) above a suitable threshold T...					
- ... i.e. one such that $T*(1+r)^N < \text{attachment point}$					
- Claims figures should include property affected, loss date, reporting date, deductibles, expenses, etc...					
- Detailed list of top losses, with description					
- Exposure data, e.g. total sum insured of properties over the same period...					
- ... or total original premium written					
- ... plus details of rate changes					
Exposure rating					
• Schedule of properties insured, with premium and sum insured/MPL for each property...					
• ... or aggregate schedule by band, such as					
Policy Limit Bands (SI, MPL)					
	Limit A	Limit B	Total Premium	Total Sum Insured/MPL	No. of Risks
	-	500,000	2,201,185	479,783,537	2603
	500,001	1,000,000	1,877,121	473,791,651	662
	1,000,001	1,500,000	1,340,030	437,375,995	353
	1,500,001	2,000,000	1,200,469	387,920,367	225
• The above must be obtained by type of property, with description of the type of property					

(b)

The credibility formula is $E(S) = \text{expected losses from experience rating,}$

$$\text{Expected_losses} = Z * \text{experience_based_average_losses} + (1-Z) * \text{exposure_based_average_losses}$$

Where

$$Z = \frac{(DS(\text{expos}))^2}{(se(E(S)))^2 + (DS(\text{expos}))^2}$$

Where $DS(\text{expos})$ is the standard deviation of the expected losses obtained by exposure rating and $se(E(S))$ is the standard error of the expected losses obtained by experience rating. Notice that in the first case $DS(\text{expos})$ is a spread over possible values of the exposure rating exercise depending on the uncertainty of what value of c is applicable, whilst $se(E(S))$ is just a result of parameter uncertainty – two conceptually very different objects!

This formula works because it balances the uncertainty on whether the market information is correct (assuming the exposure model is correct, how confident are we of our choice of c ?)

P. Parodi – Pricing in general insurance – Solutions to selected problems
with the shakiness of our experience rating calculation, which in turn depends on it being
based only on a small sample of data

(c)

The value of Z is (after dividing every number by 10,000)

$$Z = \frac{7^2}{7^2 + 3^2} = \frac{49}{58} \sim 0.845$$

The credibility estimate for the expected losses is therefore

$$EL = 0.845 * 100,000 + 0.155 * 150,000 \sim 108 \text{ (£108,000)}$$

The credibility premium can then be obtained by adding the 40% loading, yielding

$$1.4 * \text{£}108,000 \sim \text{£}151,000$$

Chapter 26

Question 1

(a)

The linear model is restricted to a function of the form $Y = a + b_1 \cdot X_1 + \dots + b_n \cdot X_n \dots$

... whereas a GLM can use any "basis" of functions each of which may include one or more factors: $Y = a + b_1 \cdot f_1(X_1, \dots, X_n) + \dots + b_N \cdot f_N(X_1, \dots, X_n)$

(note that there is an arbitrarily large number of functions f_1, \dots, f_N (possibly $N \gg n$))

The linear combination may undergo a further transformation through a link function g such that $g(Y) = a + b_1 \cdot f_1(X_1, \dots, X_n) + \dots + b_N \cdot f_N(X_1, \dots, X_n)$, further increasing the generality of the model

The noise model is not necessarily Gaussian but can be any member of the exponential family, which includes Poisson, Gamma, Gaussian, etc

(The variance may depend on the input factors) \rightarrow this is not a crucial point as heteroscedastic linear models are not that uncommon

(b)

There is no absolute rule about this but the standard model for modelling frequency is that which uses Poisson noise and a log link function.

Limiting ourselves to models which do not combine several factors, the functional form is: $\ln(Y) = a + b_1 \cdot X_1 + b_2 \cdot X_2 + \dots + b_6 \cdot X_6$ (also indicated as $\ln(Y) \sim X_1 + \dots + X_6$)

If Y is the claim count, exposure can be taken into account by using prior weights $\omega_i = \text{exposure}$.

The variance function is $V(x) = x$

(c)

Forward selection starts from the simplest model: $g(Y) = a$ (constant) and then starts adding the factors one by one...

... at each step choosing that which explains most of the residual variance (or, in other terms, for which $-\log P$ is the lowest).

The backward selection method works the other way round, starting from the full model (in our case, $g(Y) = a + b_1 \cdot X_1 + \dots + b_6 \cdot X_6$) and getting rid of all the factors one by one...

... at each step choosing that which explains most of the residual variance (or, in other terms, for which $-\log P$ is the lowest).

In both cases, you then choose the model which has the lowest AIC (Akaike Information Criterion),

$AIC = -2 \cdot \log P + 2 \cdot d$, where d = number of degrees of freedom...

... which makes sure that the finally chosen model has a good balance of fit and simplicity.

(d)

We can calculate the AIC for all the models given in the text:

	- log P	d	AIC
No factors	250	0	500
X5	165	1	332
X3, X5	145	2	294
X1, X3, X5	140	3	286
X1, X2, X3, X5	137.5	4	283
X1, X2, X3, X5, X6	137	5	284
X1, X2, X3, X4, X5, X6	136.8	6	285.6

The model with the lowest AIC is that involving the factors X1, X2, X3, X5,...

...which can be written as

$$\text{Ln}(Y) = a + b_1 \cdot X_1 + b_2 \cdot X_2 + b_3 \cdot X_3 + b_5 \cdot X_5$$

(or $\text{Ln}(Y) \sim X_1 + X_2 + X_3 + X_5$)

Question 3

(i)

Pet insurance covers the owner of a pet for such items as:

- Veterinary treatment
- Legal liability for injuries/property damage to third parties (typically, third parties exclude people living with the owner). This is normally only sold for dogs.
- Accidental death of the pet
- Pet theft or straying
- Advertisement costs/rewards for missing animals
- Owner's hospitalisation
- Holiday cancellation for reasons related to the pet

(ii)

Examples of rating factors

- Location (e.g. pet owners in London tend to pay more to reflect the higher vet fees in the capital)
- Type of pet (dog, cat, hamster...)
- Age of the animal
- Sex of the animal
- Cross breeds vs pedigree animal (pedigree's are pricier than cross breeds and more prone to certain conditions)
- Breed (some breeds have higher tendency towards certain medical conditions)
- Provenience (e.g. shelter, family, breeder...)
- Is the animal neutered/spayed?
- Existing conditions (e.g. diabetes)
- Exclusions (deductible, limits and sublimits)
- Loss history

(iii)

This goes as follows. In the following, k is the number of rating factors used.

- a. Start with the simplest model ($k=0$), i.e. $Y = a_0$.
- b. Calibrate the model, i.e. find the values of the parameters that maximise the log likelihood $\log(P)$. Call loglik the maximised value of $\log(P)$.
- c. Calculate the AIC for this model, $\text{AIC}(k) = -2 \text{loglik}(k) + 2d(k)$, where d is the number of parameters ($d=k+1$ in this case)
- d. $k \leftarrow k+1$
- e. Add one factor to the model, trying all the various factors in turn, and for each of these new models calculate the AIC – then pick the model with the lowest AIC, $\text{AIC}(k+1)$
- f. If $\text{AIC}(k+1) < \text{AIC}(k)$, the extra rating factor is accepted, and the procedure start again from d , with the addition of yet another factor
- g. If $\text{AIC}(k+1) \geq \text{AIC}(k)$, this means that the addition of a rating factor does not improve the model in a statistically significant way, and we should stop there

Chapter 27

Question 2

(a)

GLM is good at dealing with continuous variables that affect the outcome in a monotonic way (e.g. the claim size increasing monotonically with the value of the car), or with categorical variables with two values, or categorical variables with only a few values that again can be ordered.

There is a very large number of postcodes (in the order of millions), and it is not clear how they could be sorted whilst avoiding the effect of noise (small samples for a particular postcode factor).

Also, they can't easily be transformed into continuous variables, albeit two-dimensionally, as very close values (adjacent neighbourhoods) might represent very different risks

(b)

Spatial smoothing techniques allow one to reduce drastically the number of postcode-based factors that one uses...

... so that one can use the remaining values as discrete values of a categorical variable

(c)

Distance-based smoothing works by taking into account for each postcode P the information on the postcodes that are at a small enough distance from the location of P.

Eg one could replace the value of each postcode P(i) with the value $\text{average}(P(j))$ where the average is calculated over all j such that $\text{distance}(i,j) \leq d^*$...

... or more typically, one could use a weighted average, e.g. a Gaussian filter

Adjacency-based smoothing works by taking into account for each postcode the information of the neighbouring postcodes.

Eg one could replace the value of each postcode P(i) with the value $\text{average}(P(j))$ where the average is calculated over all j such that j and i are adjacent

(d)

We'll focus on the relative advantages – the relative disadvantages are the flip-side of the relative advantages, as the approaches are quite complementary as regards the points listed below.

Distance-based smoothing

(+) Easier to understand

(+) Easier to implement as it does not require distributional assumptions

(+) Can be enhanced by including extra dimensions such as urban density, income per head, etc

Adjacency-based smoothing

(+) Can take account of boundaries – eg by considering non-adjacent two postcodes across a river, regardless of the distance

(+) Prior knowledge and distributional assumptions can be incorporated in the construction of the adjacency matrix

(+) For non-weather-related perils, it deals better with changes in population density (postcodes tend to be smaller in urban regions): eg different boroughs of London may have a quite different riskiness despite their physical proximity

(e)

In distance-based smoothing, the degree of smoothing is a measure of the distance within which two location codes affect each other:

Eg if you use a 2D Gaussian filter with standard deviation = σ across both directions (symmetric filter), σ is a measure of the degree of smoothing

Or more simply, if smoothing is for you simply taking an average of everything within a radius of R , R is related to the degree of smoothing

In adjacency-based smoothing, the degree of smoothing is more difficult to define, as one simply takes everything which is adjacent – however, there is scope to increase/decrease the degree of smoothing by defining adjacency in looser or stricter terms

The degree of smoothing is designed to get rid of noise (erratic variations of risk across nearby locations):

therefore it must not be so small that the random noise is not smoothed out

but it must not be so large that useful features are also smoothed out

The optimal value can be estimated by using information criteria, which allow to establish a trade-off between noise removal and feature preservation

(f)

In time series, we also have the problem of smoothing out the signal and removing the random noise to capture the underlying dependency.

As per the case of spatial smoothing, there are several types of filters that can be used. A simple one is the so-called “moving average” by which the time series is replaced with another time series where each point is the average of k points from the original series

This is equivalent to distance-based spatial smoothing where the distance is now the length of the time series which is affected by the average

The 1D analogous of adjacency-based smoothing is more trivial – each point is simply replaced by a (possibly weighted) average of the previous point, the current point and the next point of the original series. (If the definition of “adjacency” is expanded to include more than one step backward/forward, so is the average).

Chapter 29

Question 2

(i)

The total cost of risk can be defined as follows:

Total cost of risk = Insurance premium +
 + Mean retained losses (below the EEL) +
 + Cost of capital x (Losses @ Risk acceptance level – Mean retained losses)

(ii)

One acceptable definition for the cost of capital is the following: The cost of capital is the opportunity cost for a company to set capital aside to pay for possible future losses (and therefore divest it from other projects)

(iii)

A desk quote is an insurance quote which has not been obtained directly by the insurer (i.e. “going out” in the market) but produced by a broker or a consultant on the basis of her/his experience and other elements (i.e. “at the desk”)

(iv)

We need to calculate the total cost of risk for all options and choose the option with the smallest TCOR
 The risk acceptance level is the level of probability beyond which the company doesn't care to cater for.
 In this case, RA = 100% - 1% = 99%

TCOR (Option 0 = All insured) = Premium + Mean retained + CoC x (Losses@RA – Mean retained) =
 £2,500,000 + 0 + 12% x (0 - 0) = £2,500,000

TCOR (Option 1 = EEL@ £100k) = £800k + £1.228m + 12% x (£2.000m - £1.228m) = £2.117m

TCOR (Option 1 = EEL@ £200k) = £500k + £1.602m + 12% x (£2.766m - £1.602m) = £2.242m

TCOR (Option 1 = EEL@ £500k) = £350k + £1.882m + 12% x (£3.577m - £1.882m) = £2.436m

Therefore the best option – based on the assumptions given in the text – is that Option 1.

(v)

We are comparing four options (Options 0 to 3) but for one of the options the premium is not an actual quote but a desk quote – the actual quote may turn out to be far different from expected...

... e.g. if the actual quote for Option 3 became £80k, Option 3 would become the best

The estimated TCOR is based on a loss model which will have the usual uncertainties, esp. parameter uncertainty, model uncertainty, data/assumptions uncertainty...

This, compounded with the fact that the total cost of risk for the three options with some deductible (Options 1 to 3) is very similar, makes choosing one of these options just on the basis of the TCOR unreasonable

Also, the assumptions themselves (e.g. the cost of capital, 12%, and the risk acceptance level, 99%) are likely to be the product of a not-so-certain process, which also compounds to the considerations above.