

INSTITUTE AND FACULTY OF ACTUARIES

EXAMINERS' REPORT

September 2011 examinations

Subject ST8 — General Insurance: Pricing Specialist Technical

Purpose of Examiners' Reports

The Examiners' Report is written by the Principal Examiner with the aim of helping candidates, both those who are sitting the examination for the first time and who are using past papers as a revision aid, and also those who have previously failed the subject. The Examiners are charged by Council with examining the published syllabus. Although Examiners have access to the Core Reading, which is designed to interpret the syllabus, the Examiners are not required to examine the content of Core Reading. Notwithstanding that, the questions set, and the following comments, will generally be based on Core Reading.

For numerical questions the Examiners' preferred approach to the solution is reproduced in this report. Other valid approaches are always given appropriate credit; where there is a commonly used alternative approach, this is also noted in the report. For essay-style questions, and particularly the open-ended questions in the later subjects, this report contains all the points for which the Examiners awarded marks. This is much more than a model solution – it would be impossible to write down all the points in the report in the time allowed for the question.

T J Birse
Chairman of the Board of Examiners

December 2011

General comments on Subject ST8

Subject ST8 deals with applications of general insurance pricing techniques across many different types of product. Candidates should expect the examiners to draw these applications from all parts of the syllabus in order to test as wide as possible a range of skills and, in particular, to achieve a fair balance between personal and commercial lines.

Examiners will sometimes require the use of standard general insurance and statistical techniques that are covered in earlier subjects. Candidates should ensure that they are familiar with these when preparing for the ST8 examination.

As well as pricing techniques ST8 also covers the workings and use of reinsurance products, so candidates should also expect the examiners to set questions on these aspects.

In questions with an element of calculation, different numerical answers may be obtained from those shown in these solutions depending on whether figures obtained from tables or from calculators are used in the calculations. Candidates are not penalised for this. However, candidates may be penalised where excessive rounding has been used or where insufficient working is shown. Where questions require looking up values in tables, candidates are expected to interpolate between two values if reasonable to do so, even when this is not stated in the question.

Comments on the September 2011 paper

The general performance was very similar to April 2011. Well-prepared candidates scored strongly and displayed a good understanding of the subject across the whole paper. There was no evidence of time pressure amongst the better-scoring candidates.

There was a good spread of marks amongst candidates on most questions, but Q4 and Q5 in particular produced relatively low scores. Apart from those, Q3(ii), Q7(iii)-(iv), Q8(v-viii) and Q9(iv) appeared to be the most difficult and tended to discriminate the better candidates.

The comments that follow the questions concentrate on areas where candidates could have improved their performance. Candidates approaching the subject for the first time are advised to concentrate their revision in these areas.

- 1 (i)
- To limit exposure to risk or to spread risk
(no credit for simply saying it increases diversification)
 - E.g., single risks, aggregations of single risks, accumulations, multi-class losses
 - To avoid single large losses e.g. liability claims
 - Reinsurance can increase the opportunity for an insurer to make a profit and plan its business more accurately
(no credit for simply saying it increases profits)
 - To smooth results
 - To improve solvency margins or reduce the required solvency margin
 - To increase an insurer's capacity to accept risk
 - To gain expertise when developing new markets / products
 - To participate in reciprocal arrangements
 - To gain financial assistance
 - e.g. against new business strain
 - For legislative reasons
 - e.g. a compulsory terror pool
 - Could offer tax advantages

This bookwork part was generally well-answered.

- (ii) EML = £21.6m
R = £3m
So the number of lines ceded = $(21.6 / 3) - 1 = 6.2$ lines

Therefore all claims will be split in the proportion 6.2 : 1

Claim is £24.5m (it doesn't matter that this is more than the EML)
...and is also split in the proportion 6.2 : 1

Reinsurer pays $(6.2 / 7.2) * 24.5m = £21.097m$

Some candidates did not realise that the claim would be split in the treaty proportions even though it is above the EML. A few candidates appeared not to understand the operation of a surplus treaty properly.

- 2 (i) Inflate the claims to current day values
Develop to ultimate (add IBNER)
Use different trend rates for expenses and indemnity
Limit inflated, developed indemnity to \$1m
(must be clear that capping applies after inflation & development)
Consider whether expenses require adjustment as a result of the new indemnity limit and adjust if so.

Very few candidates mentioned using different adjustments for expenses and indemnity elements. A common error was to state that indemnity amounts should be limited before claim amounts were inflated.

- (ii) Many claims over \$1m will reduce
Not all, though, as there will always be expenses
Claims with an indemnity amount under \$1m will remain the same
So there will be a flattening of the curve over \$1m
The answer is C

Most candidates recognised that the correct curve was C but few were able to explain clearly why.

- 3** (i) Main advantage – they can allow for the “soft” factors that would otherwise be unquantifiable.

e.g. subtle changes in terms and conditions, risk management changes, or other sensible non-quantifiable factor

Quicker/cheaper to compute

Makes use of underwriter's experience & knowledge

Main disadvantage – very subjective (depends on underwriter asked)

Difficult to ensure consistency over time.

Difficult to assess across companies and classes.

Difficult to verify or quantify in detail analytically.

Easy to manipulate – underwriter may have a vested interest or bias.

No audit trail of calculations (for a regulator or auditor).

There may be confusion as to whether a change in the premium is linked to the level of risk or is a rate change.

No credit for saying this is useful as a check on other methods because question says “relying”.

This bookwork part was generally well-answered.

- (ii) As-if prem(09) =

$$\text{Prem}(09) \times \frac{\text{ILF@Limit}(10) - \text{ILF@Attach}(10)}{\text{ILF@Limit}(09) - \text{ILF@Attach}(09)} \times \frac{\text{Share}(10)}{\text{Share}(09)} \times \frac{\text{Exposure}(10)}{\text{Exposure}(09)}$$

$$= \text{Prem}(09) \times \frac{(2.900 - 2.300)}{(2.750 - 1.000)} \times \frac{23}{20} \times \frac{1,500,000}{1,000,000}$$

$$= £4,900 \times \frac{0.6}{1.75} \times \frac{23}{20} \times 1.5$$

$$= \text{£}2,898$$

Hence the change in premium rate from 2009 to 2010 is:

$$\begin{aligned} &= \frac{\text{Prem rate}(10)}{\text{Prem rate}(09)} = \frac{\text{Prem}(10)}{\text{As-if prem}(09)} \\ &= \frac{2,700}{2,898} \\ &= 0.931677 \end{aligned}$$

The true change in rate is -6.83%

Alternative calculation method:

$$\text{Total premium for 2009} = 4900 / 20\% = 24500$$

$$\text{Total premium for 2010} = 2700 / 23\% = 11739$$

$$\text{ILF for 2009 cover} = 2.75 - 1 = 1.75$$

$$\text{ILF for 2010 cover} = 2.9 - 2.3 = 0.6$$

$$\text{Premium per unit cover 2009} = 24500 / 1.75 = 14000$$

$$\text{Premium per unit cover 2009} = 11739 / 0.6 = 19565$$

$$\text{Premium per unit turnover \& cover 2009} = 14000 / 1,000,000 = 0.014$$

$$\text{Premium per unit turnover \& cover 2010} = 19565 / 1,500,000 = 0.013$$

$$\text{Rate change from 2009 to 2010} = 0.014 / 0.013 = 0.931677$$

So the true rate change is minus 6.83%

Candidates came up with a wide range of approaches and answers to this part, which was surprising given that the first method above appears in Core Reading. Many candidates picked up some credit for an alternative method but then lost their way. Despite the question making it very clear what the extent of cover was, many candidates failed to add the policy limit to the excess to find the upper ILF and ended up interpolating between two values from the table for 2009.

4 Consider data availability Consider data quality

More years gives more credibility

We want sufficient data to smooth out random volatilities

EL can have very volatile claims experience

We should also consider the **completeness** of the claims data

We need to go sufficiently far back to get:

- The full range of large losses

- Catastrophe type losses
- Experience in some of the more low-frequency rating cells

More heterogeneity in the book means more data required and hence more years

Consider the complexity of the model to be built (e.g. frequency/severity needs more data than aggregate)

Influences on this include:

- The size and age of the company
- Quality and integrity of systems and processes
- Availability of data from external sources

Older years will be less **relevant** to current experience

Examples of losing relevance ($\frac{1}{4}$ each – max 1):

- big changes in risk or mix of business;
- change in underwriting practice
- changes in claims handling practice
- change in the legal environment for claims
- change in propensity to claims;
- different cover,
- different types of claim;
- more difficult to inflate accurately.

We need a certain number of years to identify **trends**

More recent experience is more **uncertain** (unsettled and non-reported claims)
Therefore we may drop more recent years

Especially true in a long-tailed class like EL

This question was quite straightforward and mainly well-covered in Core Reading but generally produced lower scores than the other questions. Candidates tended to regurgitate a limited number of facts about the class of business, rather than trying to answer the question. Most candidates were able to give a range of examples of the lack of relevance of data from older years, but frequently omitted points related to going sufficiently far back to get the full range of losses, cat losses and experience in the low-frequency cells.

5

(i)

- Windstorm/Typhoon/Hurricane/Tropical storm
- Earthquake
- Tornado
- Hail
- Winter storm/Freeze/Snow
- Flood
- Tsunami/tidal wave
- Wildfire/forest fire (or other widespread fire)

- Terrorism
- Pollution (often excluded)
- Nuclear (often excluded)

This part was generally well-answered but some candidates listed perils that would not be covered in catastrophe reinsurance, such as theft.

(ii) Reasons

- As with any model, there may be approximations and lack of fit.
- Equally, if the output of the model were the same as actual experience over a long period, this would suggest over-fitting/lack of predictive power.
- The model might be out of date.
- An incomplete event set
 - We may be missing extreme events
 - Some of our exposures may not be adequately covered by the hurricane paths
 - We may be generating from past experience, which does not account for claims trends
 - e.g. in the Gulf of Mexico, due to global warming; economic recession (or other suitable example)
- Errors in the hazard model
 - e.g. wind speed too low, diameter not wide enough
- May under/over estimate potential losses in the vulnerability model
 - For example demand surge may not be sufficiently modelled; or construction types respond unexpectedly; or flood defences perform differently; or other suitable example
- May not model some of the exposures
 - e.g. unusual occupancy or construction type
- Wrong perils switched on in the model
- Coverage not correctly modelled, e.g. flooding excluded
- User input error or mistakes in exposure sheet
- Exposures incomplete
- The model will have been based on expectations of the exposures and mix of business that the reinsurer would take on, but in reality this may turn out to be different.

No credit for mentioning volatility of underlying claims experience.

Candidates did not generally score well on this part because they failed to structure their answer and give a wide enough range of points. The stronger candidates broke their answer down into sections relating to the five sections of a catastrophe model and found that this helped to generate ideas. A significant number of candidates forgot that the three events stated in the question were only an extract of the output, and made comments that it was not appropriate to build a catastrophe model with only three historic events.

- (iii) Whether each event is covered under the ILW
The industry loss for each event

- (iv)
- For each loss (col (2) above) calculate the recoveries if the ILW were triggered (new Col (5))
 - For each event use the industry loss to see if the ILW is triggered
 - Multiply the recovery (Col 5) by their frequency and sum up all the triggered rows

It was evident from answers to (iii) and (iv) that most candidates had a sketchy understanding of the operation of an ILW. In (iii), many candidates gave "industry loss" as an answer but did not say "for each event" and failed to make the first point. Part (iv) was very poorly answered, with most candidates failing to be sufficiently precise to convey understanding.

- 6** (i) incorrect advice
error in calculation or report

A surprising number of candidates failed to notice that the question stated the insured is an actuarial consultancy, and suggested perils relating to medical malpractice.

- (ii)
- Limit
 - Deductible
 - No. of actuaries/employees
 - Payroll
 - Turnover
 - Location of HQ
 - Territory of practice
 - Claims experience
 - Area of practice e.g. GI, Life, Pensions
 - Type of work e.g. opinions, reserving, M&A
 - Type of client (eg government)
 - Additional coverages e.g. public liability (PL), extra contractual obligations (ECO) & excess of policy limits (XPL)
 - Exclusions e.g. punitive damages

This part was generally well-answered.

- (iii) Claims-made covers all claims first notified within the policy period irrespective of when the event occurred

provided that this is after the retroactive date

On losses-occurring cover claims event must have occurred during the policy period

A claims made policy can be taken out to cover events that may already have occurred

Claims-made basis may give unsatisfactory cover for future claimants where the tortfeasor (i.e. defendant) may cease to exist or cannot obtain cover in the future

Most candidates gave the definition of claims made and losses occurring but did not give enough further details to gain full marks.

- (iv) No. of actuaries
 No. of partners
 No. of billable hours

Many candidates failed to score the full mark because they gave answers that were not suitable measures of exposure or would have to be adjusted for inflation.

(v)

<i>Policy Year</i>	<i>CM Exposure</i>
2009	85.175
2010	87.375
2011	90.475

= 95*0.25+90*0.5+86.9*0.25

This part was generally well-answered.

(vi)

<i>Policy Year</i>	<i>CM Exposure</i>
2009	21.725
2010	65.95
2011	90.475

= 86.9*.25
 = 90*0.25+86.9*0.5
 = 95*0.25+90*0.5+86.9*0.25

Assumptions:

- No difference in value of claims with longer reporting delay
- Uniform incidence of occurrence of risk(/claims) throughout the exposure year

Few candidates gave the required assumptions in this part.

7 (i) Quota share reinsurance:

- Proportional reinsurance
- Claims and premiums shared by an agreed proportion for each risk
- Proportion same for each risk
- Administered by treaty
- May involve an overriding commission (additional commission payable from reinsurer to insurer as a contribution to expenses and profit)
- May also involve a profit commission

- The treaty may specify a limit on the amount of business that may be ceded
- Usually written on a policies incepting basis

Many candidates wasted time by giving applications of the cover, for which there was no credit.

(ii)

Year of Account	On-Level Ultimate Loss Ratio	
2008	55.13%	$59\% / ((1-2\%)*(1+5%)*(1+4\%))$
2009	46.70%	$51\% / ((1+5%)*(1+4\%))$
2010	40.95%	$43\% / (1+5\%)$

A common error in this part was multiplying ULR by the required adjustments instead of dividing.

(iii) Points:

- Easy to understand/explain
- Has the benefit of being based on actual experience
- However, past experience may not be a good guide to the future
- Profit = premium \times (1 – exp) – ult. losses
- So to get zero profit we have:
- Loss ratio = 1 – exp = 70%
- i.e. profit is paid out on loss ratios < 70%
- So none of the historic loss ratios would pay out a commission
- This would give a PC load of zero
- This is unrealistic as there must be some chance of paying out
- We only have 11 years of data here. If we had many more we would have the variability to trigger the PC

(iv) Points:

- Probably a good fit to the data, since the mean and variance are sample statistics
- Easy to calculate/apply
- This approach would give a wide variety of results triggering the profit commission, giving a more realistic approach to a long term average
- Aggregate claims distribution tend to be skewed. A normal distribution does not reflect this
- Specifically a normal distribution can go negative and typically does not have a long tail
- However this distribution is very tight and in reality going negative is very extreme
- In addition for PC we're only interested in the distribution below 70% so the tail doesn't matter

- We may wish to give more weight to more recent experience, which this approach does not do
- The past experience might not be a good guide to the future, so the parameters of the distribution may be inappropriate
- In particular, the fact that C has requested a profit commission for the first time suggests a different approach to underwriting in the future

Many candidates misunderstood the operation of the profit commission in parts (iii) and (iv) despite the clear description in the question. Some forgot that it only applied to a single contract and others thought that it was determined using historic experience rather than 2011 performance. In (iii) many candidates concentrated on the long term average loss ratio instead of noting that none of the individual years would have resulted in a profit share being paid. Candidates should note that it is often important to observe the features of the data given in questions.

(v) $Z = 1.778 \quad (.86 - .7) / .09$

Interpolated	x_i	$\Phi(x_i)$	Wt_i
	1.7700	0.96164	22%
	1.7800	0.96246	78%

$F(x)$ 96.23%
 prob 3.77%

Many candidates rounded to a value of 1.78 instead of interpolating between 1.77 and 1.78, thereby throwing away easy marks.

- (vi) No effect.
 Claims and Premiums will be scaled by the same amount.

Most candidates understood that it would have no effect but many were unable to explain clearly why.

8 (i)

Trend Factor	Trended Claims (£)						Number
1.4071	9,850	19,699	49,249	35,178	-	-	-
1.3401	95,147	120,609	45,563	-	-	-	1
1.2763	104,655	70,195	236,112	66,367	-	-	2
1.2155	125,197	29,172	4,862	179,895	269,842	-	3
1.1576	19,680	226,895	360,021	6,946	670,265	-	3
1.1025	449,820	110,250	67,253	45,203	759,623	429,975	4
1.0500	77,700	134,400	242,550	229,950	67,200	55,650	3

An alternative and quicker approach (still for full marks) would be to deflate the deductible:

<i>Deflator</i>	<i>Deductible</i>	<i>Number</i>
0.7107	71,068	-
0.7462	74,622	1
0.7835	78,353	2
0.8227	82,270	3
0.8638	86,384	3
0.9070	90,703	4
0.9524	95,238	3

(ii)

<i>Policy Year</i>	<i>% Reported</i>	<i>Ultimate Number</i>
2004	100%	-
2005	100%	1.000
2006	100%	2.000
2007	95%	3.158
2008	90%	3.333
2009	80%	5.000
2010	55%	5.455

(iii) Identify an upwards trend

There are a number of different selections we can make. Sensible selection
Corresponding explanation

e.g.

- 6–7 frequency is increasing dramatically – extrapolated this
- 5–6 select the most recent experience (seems to be higher than historical)
- 3–5 average over last 4 years (last 2 years may be abnormally poor and '10 is uncertain)

Allow claim frequency per ship instead of per policy

Parts (i), (ii) and (iii) were all generally well-answered

(iv) The upward trend in claims may be partly/completely caused by the increase in exposure (ship numbers)

Therefore there may not be a upwards trend of claims/exposure

We need to analyse ultimate frequency (or ultimate aggregate claims) / exposure to identify any trends

This requires obtaining all data on ship numbers or perhaps an alternative exposure measure, such as weight or total insured value

May want to enquire why the increase e.g. new types of business, new locations etc.

Many candidates suggested that the number of claims should be inflated rather than re-examining claim frequency.

(v)

- Different deductible levels may affect the reporting pattern because the insured may report smaller claims at a different speed
- This may be applied for a different type of coverage e.g. liability claims excluded
- The mix of business may be different (eg type of ship)
- Insured may transact same type of business but in a territory with different reporting speeds
- The basis for estimating outstanding claims may be different.
- Insured may have different risk attitude or complaints processes, affecting the reporting speed

This part was answered well by the majority of candidates.

(vi) Poisson
Negative binomial

Most candidates scored full marks for this part.

(vii) The individual paid plus outstanding claims estimates may need further development (IBNER adjustment)
Ideally only on open claims
The resulting distribution may also need adjustment to allow for IBNR

Many candidates did not make the points in enough detail to score the full mark. Simply saying that the claims needed to be developed was not sufficient.

(viii)

- Quick and easy approach
- Would expect claim numbers to increase proportionally to exposure
- The new shipping company may have higher or lower expected claims costs than the current company because of:
 - Different types of ships /types of cargo / age of ships / quality of ships
 - Different experience of crew
 - Different territories/legal jurisdiction;
 - Differences in risk management.

- We should be especially wary of the different frequency/severity dynamics here. The 25% method probably assumes frequency per ship and severity per claim are the same for the two companies. If frequency or severity of claims is different, the layered results would be difficult to estimate without remodelling. Conversely there may be a large increase in small claims (under the deductible) but not much on larger ones.
- Ideally we would want all GU claims and historic ship numbers for the new entity
- ...and model the combined company
- Data might not be easily available/reliable
- The expected loss will not be exactly proportional to ship numbers due to the aggregate limit
- Alternatively, the aggregate limit may not be appropriate and may need to be increased
- There may be aggregations that cause diversification to worsen (alternatively, acquisition might increase diversification)
- Larger shipping companies may get a size discount due to better risk management from having specialist risk management departments (made possible due to scale)
- ...however probably too early for this.
- However if acquisition unpopular amongst staff claims may increase e.g. safety officers all leave
- Any additional reinsurance costs might not be proportional (could be higher or lower)
- Any other additional expenses might not be proportional (could be higher or lower)

This part was not generally well-answered. The better-scoring candidates gave a wide range of points and clear explanations.

9 (i) $\mathbf{Y} = g^{-1}(\mathbf{X}\boldsymbol{\beta} + \boldsymbol{\xi}) + \boldsymbol{\varepsilon}$
(or $Y_i = g^{-1}(\sum X_{ij} \beta_j + \xi_i) + \varepsilon_i$)

where:

\mathbf{Y} is the response vector

$g(\)$ is the link function

\mathbf{X} is the design matrix of factors

$\boldsymbol{\beta}$ is the vector of parameters to be estimated

$\boldsymbol{\xi}$ is a vector of offsets or known effects

$\boldsymbol{\varepsilon}$ is the error term appropriate to \mathbf{Y}

This part was generally well-answered.

- (ii) A categorical factor is a factor for which the values of each level are distinct and often cannot be given any natural ordering or score.

A factor which is not categorical is one that takes a naturally ordered value.

This part was generally well-answered. Some candidates gave an example that applied to motor insurance, which was surprising in a question about household contents.

- (iii) The scaled deviance and AIC are both statistics used to assess which of two models is the better fit.

If the two models are nested and the scale parameter is known, a χ^2 -test is used.

The change in scaled deviance ($D_1^* - D_2^*$) is compared with a χ^2 -distribution with $(df_1 - df_2)$ degrees of freedom.

If the two models are nested but the scale parameter is not known, an F -test is used.

The statistic $(D_1 - D_2) / ((df_1 - df_2)(D_2 / df_2))$ is compared with an F -distribution with $(df_1 - df_2, df_2)$ degrees of freedom.

If the two models are not nested, the AIC is used
where $AIC = -2 * \text{loglikelihood} + 2 * \text{number of parameters}$

The model with the lower AIC is said to have the better fit.

This part was generally well-answered. However, few candidates mentioned the F -test.

- (iv) Models 1A and 1B are subsets of Model 0 so comparisons with Model 0 should use the scaled deviance.

The difference in scaled deviance for Model 1A = 1.1
and the difference in degrees of freedom = 1.

The upper 5% point of χ^2_1 is 3.841 and $1.1 < 3.841$ (equal credit for calculating a p -value of around 0.3)

So Model 1A is not significantly different from Model 0 and so we would conclude that the “occupied during the day” indicator is not a significant factor.

We would reject Model 0 in favour of Model 1A.

The difference in scaled deviance for Model 1B = 11.3
and the difference in degrees of freedom = 6.

The upper 5% point of χ^2_6 is 12.59 and $11.3 < 12.59$ (equal credit for calculating a p-value of around 0.08)

So Model 1B is not significantly different from Model 0 and so we would conclude that the “property type” indicator is not a significant factor.

We would reject Model 0 in favour of Model 1B.

Models 1A and 1B are not nested but a further comparison between these can still be done using AIC.

The AIC for Model 1A is lower than that for 1B, suggesting that Model 1A is preferred.

Dropping property type in 1B does not appear as successful as dropping “occupied during the day” in 1A.

AIC includes a penalty for the number of parameters. Although dropping property type removes a larger number of parameters from the model (six as opposed to just one), it has more influence on the fit of the model, so the AIC is poorer for 1B.

It is perhaps surprising that these rating factors are not more significant, which suggests that other factors are acting as proxies, or that there may be an error in the data or model.

A large number of candidates used the wrong number of degrees of freedom for the χ^2 tests. Even of those who calculated the test statistic correctly a considerable number drew the wrong conclusion from the results of the tests.

END OF EXAMINERS' REPORT