Clark

Outline

I. Introduction

◇ Objectives in creating a formal model of loss reserving:

- Describe loss emergence in simple mathematical terms as a guide to selecting amounts for carried reserves
- Provide a means of estimating the range of possible outcomes around the "expected" reserve

◇ A statistical loss reserving model has two key elements:

- The expected amount of loss to emerge in some time period
- The distribution of actual emergence around the expected value

II. Expected Loss Emergence

◇ Model will estimate the expected amount of loss to emerge based on:

- An estimate of the ultimate loss by year
- An estimate of the pattern of loss emergence

◇ Let $G(x) = 1/LDF_x$ be the cumulative % of loss reported (or paid) as of time x, where x represents the time (in months) from the "average" accident date to the evaluation date

◇ Assume that the loss emergence pattern is described by one of the following curves with scale θ and shape ω

- Loglogistic

$$\boxed{\begin{aligned} G(x|\omega,\theta) &= \frac{x^\omega}{x^\omega + \theta^\omega} \\ LDF_x &= 1 + \theta^\omega \cdot x^{-\omega} \end{aligned}}$$

- Weibull

$$\boxed{G(x|\omega,\theta) = 1 - \exp(-(x/\theta)^\omega)}$$

◇ With these curves, we assume a strictly increasing pattern. If there is real expected negative development (salvage recoveries), different models should be used

◇ Advantages to using parameterized curves to describe the emergence pattern:

- Estimation is simple since we only have to estimate two parameters
- We can use data that is not from a triangle with evenly spaced evaluation data – such as the case in which the latest diagonal is only nine months from the second latest diagonal
- The final pattern is smooth and does not follow random movements in the historical age-to-age factors

◇ In order to estimate the loss emergence amount, we require an estimate of the ultimate loss by AY. There are two methods described in the paper:

- LDF method – assumes the loss amount in each AY is independent from all other years (this is the standard chain-ladder method)
- Cape Cod method – assumes that there is a known relationship between expected ultimate losses across accident years, where the relationship is identified by an exposure base (on-level premium, sales, payroll, etc.)

◇ Let $\mu_{AY;x,y}$ = expected incremental loss dollars in accident year AY between ages x and y

◇ Combining the loss emergence pattern with the estimate of the ultimate loss by year, we obtain the following for each method:

- LDF method

$$\boxed{\mu_{AY;x,y} = ULT_{AY} \cdot [G(y|\omega,\theta) - G(x|\omega,\theta)]}$$

- Cape Cod method

$$\boxed{\mu_{AY;x,y} = \text{Premium}_{AY} \cdot ELR \cdot [G(y|\omega,\theta) - G(x|\omega,\theta)]}$$

◇ In general, the Cape Cod method is preferred since data is summarized into a loss triangle with relatively few data points. Since the LDF method requires an estimation of a number of parameters (one for each AY ultimate loss, as well as θ and ω), it tends to be over-parameterized when few data points exist

◇ Due to the additional information given by the exposure base (as well as fewer parameters), the Cape Cod method has a smaller parameter variance. The process variance can be higher or lower than the LDF method. In general, the Cape Cod method produces a lower total variance than the LDF method

III. The Distribution of Actual Loss Emergence and Maximum Likelihood

◇ The variance of the actual loss emergence can be estimated in two pieces: **process variance** (the "random" amount) and **parameter variance** (the uncertainty in the estimator, also known as the estimation error)

◇ **Process variance**

- Assume that the loss in any period has a constant ratio of variance/mean:

$$\boxed{\frac{\text{Variance}}{\text{Mean}} = \sigma^2 \approx \frac{1}{n-p} \sum_{AY,t}^{n} \frac{(c_{AY,t} - \mu_{AY,t})^2}{\mu_{AY,t}}}$$

  where n = # of data points, p = # of parameters, $c_{AY,t}$ = actual incremental loss emergence and $\mu_{AY,t}$ = expected incremental loss emergence

- For estimating the parameters of our model, let's assume that the actual loss emergence "c" follows an over-dispersed Poisson distribution with scaling factor σ^2

- Assuming λ represents the mean of a standard Poisson random variable, the mean and variance of an over-dispersed Poisson are as follows:

  ◇ $E[c] = \lambda\sigma^2 = \mu$

  ◇ $Var(c) = \lambda\sigma^4 = \mu\sigma^2$

- Key advantages of using the over-dispersed Poisson distribution:

  ◇ Inclusion of scaling factors allows us to match the first and second moments of any distribution, allowing high flexibility

  ◇ Maximum likelihood estimation produces the LDF and Cape Cod estimates of ultimate losses, so the results can be presented in a familiar format

◇ **The likelihood function**

- For an over-dispersed Poisson distribution, the $\Pr(c) = \frac{\lambda^{c/\sigma^2} e^{-\lambda}}{(c/\sigma^2)!}$

- Likelihood $= \prod_i \Pr(c_i) = \prod_i \frac{\lambda_i^{c_i/\sigma^2} e^{-\lambda_i}}{(c_i/\sigma^2)!} = \prod_i \frac{(\mu_i/\sigma^2)^{c_i/\sigma^2} e^{-(\mu_i/\sigma^2)}}{(c_i/\sigma^2)!}$

- After taking the log of the likelihood function above, we obtain the loglikelihood, l, which we need to maximize:

$$\boxed{l = \sum_i c_i \cdot \ln(\mu_i) - \mu_i}$$

- Before applying this loglikelihood formula to our two methods, let's define a few things:

  ◇ $c_{i,t}$ = actual loss in AY i, development period t

  ◇ P_i = premium for AY i

  ◇ x_{t-1} = beginning age for development period t

  ◇ x_t = ending age for development period t

- LDF method

- ◇ Taking the derivative of l and setting it equal to zero yields the following MLE estimate for ULT_i:

$$ULT_i = \frac{\sum_t c_{i,t}}{\sum_t [G(x_t) - G(x_{t-1})]}$$

- ◇ The MLE estimate for each ULT_i is equivalent to the "LDF Ultimate"

- Cape Cod method

  - ◇ Taking the derivative of l and setting it equal to zero yields the following MLE estimate for the ELR:

$$ELR = \frac{\sum_{i,t} c_{i,t}}{\sum_{i,t} P_i \cdot [G(x_t) - G(x_{t-1})]}$$

  - ◇ The MLE estimate for the ELR is equivalent to the "Cape Cod" Ultimate

- An **advantage** of the maximum loglikelihood function is that it works in the presence of negative or zero incremental losses (since we never actually take the log of $c_{i,t}$)

◇ **Parameter variance**

- We need the covariance matrix (inverse of the information matrix) to calculate the parameter variance
- Due to the complexity involved (it would be downright impossible for the LDF method), I don't expect you will need to calculate the parameter variance on the exam

◇ **Variance of the reserves**

- As usual, in order to calculate the variance of an estimate of loss reserves R, we need the process variance and parameter variance:
  - ◇ Process Variance of $R = \sigma^2 \cdot \sum \mu_{AY;x,y}$
  - ◇ Parameter Variance of R = too complicated for the exam

IV. Key Assumptions of this Model

◇ Assumption 1: Incremental losses are independent and identically distributed (iid)

- "Independence" means that one period does not affect the surrounding periods
  - ◇ Can be tested using residual analysis
  - ◇ Positive correlation could exist if all periods are equally impacted by a change in loss inflation

  - ◇ Negative correlation could exist if a large settlement in one period replaces a stream of payments in later periods
- "Identically distributed" assumes that the emergence pattern is the same for all accident years, which is clearly over-simplified
  - ◇ Different risks and a different mix of business would have been written in each historical period, each subject to different claims handling and settlement practices

◇ Assumption 2: The variance/mean scale parameter σ^2 is fixed and known

- Technically, σ^2 should be estimated simultaneously with the other model parameters, with the variance around its estimate included in the covariance matrix
- However, doing so results in messy mathematics. For convenience and simplicity, we assume that σ^2 is fixed and known

◇ Assumption 3: Variance estimates are based on an approximation to the Rao-Cramer lower bound

- The estimate of variance based on the information matrix is only exact when we are using linear functions
- Since our model is non-linear, the variance estimate is a Rao-Cramer lower bound (i.e. the variance estimate is as low as it possibly can be)

V. A Practical Example

◇ *In the paper, Clark applies his methodology to 10 x 10 triangle. To simplify things, we will be studying a 5 x 5 triangle. In general, this example will focus on estimating the reserves using the LDF and Cape Cod methods. For the more detailed calculations (such as determining model parameters or calculating residuals), see the Clark Example excel spreadsheet within the online course.*

◇ Before diving into the example, let's briefly discuss growth curve extrapolation:

- The growth curve extrapolates reported losses to ultimate
- For curves with "heavy" tails (such as loglogistic), it may be necessary to truncate the LDF at a finite point in time to reduce reliance on the extrapolation
- An alternative to truncating the tail factor is using a growth curve with a "lighter" tail (such as Weibull)

◇ **LDF method**

- Assume that expected loss emergence is described by a loglogistic curve. In addition, assume that the curve should be truncated at 120 months
- Given the following cumulative losses and parameters:

	Cumulative Losses ($)				
AY	12	24	36	48	60
2010	500	1500	2250	2590	2720
2011	550	1700	2400	2725	
2012	450	1200	2000		
2013	600	1750			
2014	575				

Parameters	
θ	21.4675
ω	1.477251
σ^2	59.9876

- Create the following table to estimate the reserves:

AY	Losses at 12/31/14	Age at 12/31/14	Avg. Age (x)	Growth Function	Fitted LDF	Trunc. LDF	Estimated Reserves	Estimated Ultimate
Trunc.		120	114	0.922	1.0846	1.0000		
2010	2720	60	54	0.796	1.2563	1.1583	430.576	3150.576
2011	2725	48	42	0.729	1.3717	1.2647	721.308	3446.308
2012	2000	36	30	0.621	1.6103	1.4847	969.400	2969.400
2013	1750	24	18	0.435	2.2989	2.1195	1959.125	3709.125
2014	575	12	6	0.132	7.5758	6.9848	3441.260	4016.260
Total							7521.669	17291.669

- Here are the 2013 calculations for the table above:
  - ◇ Avg. age $= 18 = 24 - 6$
  - ◇ Growth function $= \frac{x^\omega}{x^\omega+\theta^\omega} = \frac{18^{1.477251}}{18^{1.477251}+21.4675^{1.477251}} = 0.435$
  - ◇ Fitted LDF $= \frac{1}{0.435} = 2.2989$
  - ◇ Truncated LDF $= \frac{0.922}{0.435} = 2.1195$
  - ◇ Estimated reserves $= 1750(2.1195 - 1) = 1959.125$
  - ◇ Estimated ultimate $= 1750 + 1959.125 = 3709.125$
- To calculate the process standard deviations of the reserves for each accident year, we multiply the scale parameter σ^2 by the estimated reserves and take the square root. Thus, we have the following:

AY	Estimated Reserves	Process SD
2010	430.576	160.715
2011	721.308	208.013
2012	969.400	241.147
2013	1959.125	$342.817 = \sqrt{59.9876(1959.125)}$
2014	3441.260	454.349
Total	7521.669	671.719

◇ **CC method**

- Assume that expected loss emergence is described by a Loglogistic curve. In addition, assume that the curve should be truncated at 120 months
- Given the following cumulative loss and parameters:

	Cumulative Losses ($)				
AY	12	24	36	48	60
10	500	1500	2250	2590	2720
11	550	1700	2400	2725	
12	450	1200	2000		
13	600	1750			
14	575				

Parameters	
θ	22.3671
ω	1.441024
σ^2	50.0730

- Create the following table to calculate the ELR (**note that the ELR is calculated before truncation to remain algebraically consistent with how the LDF method works**):

AY	On-Level Premium	Losses at 12/31/14	Age at 12/31/14	Avg. Age (x)	Growth Function	Premium × Growth
2010	5000	2720	60	54	0.781	3905.00
2011	5200	2725	48	42	0.713	3707.60
2012	5400	2000	36	30	0.604	3261.60
2013	5600	1750	24	18	0.422	2363.20
2014	5800	575	12	6	0.131	759.80

- Here are the 2013 calculations for the table above:
  - ◇ Average age $= 18 = 24 - 6$
  - ◇ Growth function $= \frac{x^\omega}{x^\omega + \theta^\omega} = \frac{18^{1.441024}}{18^{1.441024} + 22.3671^{1.441024}} = 0.422$
  - ◇ Premium × growth $= 5600(0.422) = 2363.20$
- The expected loss ratio is $\frac{2720+2725+2000+1750+575}{3905+3707.60+3261.60+2363.20+759.80} = 0.698$
- Assuming a truncation point of 120 months, estimate the reserves:

AY	On-Level Premium	Age at 12/31/14	Average Age (x)	Growth Function	0.913 − Growth	Expected Losses	Estimated Reserves
Trunc.		120	114	0.913	0.000		
2010	5000	60	54	0.781	0.132	3490.00	460.680
2011	5200	48	42	0.713	0.200	3629.60	725.920
2012	5400	36	30	0.604	0.309	3769.20	1164.683
2013	5600	24	18	0.422	0.491	3908.80	1919.221
2014	5800	12	6	0.131	0.782	4048.40	3165.849
Total							7436.353

- For 2013, the expected losses are $3908.8 = 5600(0.698)$ and the estimated reserves are $1919.221 = 3908.8(0.491)$

- Here are the process standard deviations:

AY	Estimated Reserves	Process SD
2010	460.680	151.880
2011	725.920	190.654
2012	1164.683	241.494
2013	1919.221	$310.002 = \sqrt{50.0730(1919.221)}$
2014	3165.849	398.150
Total	7436.353	610.213

◇ Residuals

- The scale factor σ^2 is useful for a review of the model residuals, $r_{AY;x,y}$:

$$r_{AY;x,y} = \frac{c_{AY;x,y} - \hat{\mu}_{AY;x,y}}{\sqrt{\sigma^2 \cdot \hat{\mu}_{AY;x,y}}}$$

- We plot the residuals against a number of things to test model assumptions:
  - ◇ Increment age (i.e. AY age)
  - ◇ Expected loss in each increment - useful for testing if variance/mean ratio is constant
  - ◇ Accident year
  - ◇ Calendar year - to test diagonal effects
- In all of the cases above, we want the residuals to be randomly scattered around the zero line
- Here is an example of a residual graph for the LDF method shown above:

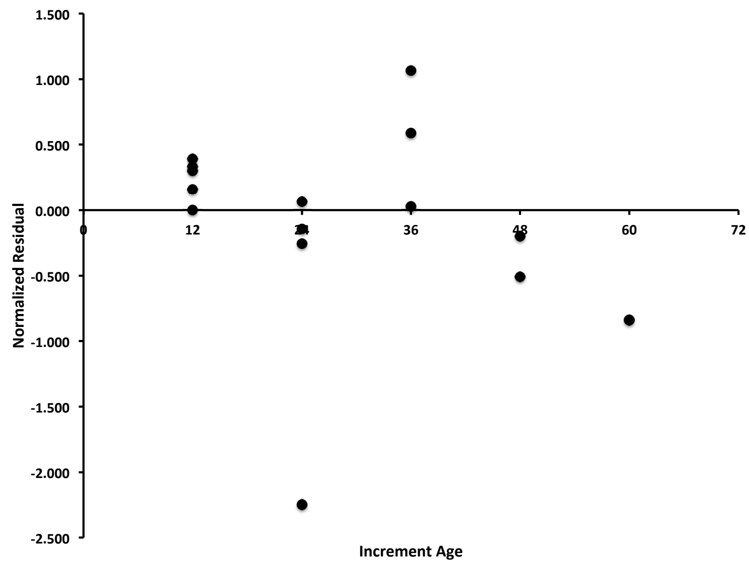

- In this case, the residuals do NOT appear to be randomly scattered around the zero line. Thus, we conclude that the model assumptions are invalid

◇ Testing the constant ELR assumption in the Cape Cod model

- Graph the ultimate loss ratios AY, where the ultimate loss ratio is equal to the reported losses divided by the used-up premium; this is equivalent to the loss ratios from the LDF method
- If an increasing or decreasing pattern exists, this assumption may not hold
- As an example, consider the following:

AY	On-Level Premium	Reported Losses	Growth Function	Used-Up Premium	Ultimate Loss Ratio
2014	1000	600	0.623	623	$\frac{600}{623} = 0.963$
2015	1200	500	0.472	566.4	0.883
2016	1400	180	0.178	249.2	0.722

- In this case, there is an obvious decreasing pattern in the ultimate loss ratios. Thus, the constant ELR assumption does not appear to hold

◇ Other calculations possible with this model

- Variance of the prospective losses
  - ◇ Uses the Cape Cod method
  - ◇ If we have an estimate of future year premium, we can easily calculate the estimate of expected loss (which in this case would be the estimated reserves) because we already have the maximum likelihood estimate of the ELR

- ◇ The process variance is calculated as usual
- ◇ For example, if the maximum likelihood estimate of the ELR is 0.75 and next year's planned premium is \$6M, then the prospective losses for next year are \$6M(0.75) = \$4.5M. Given $\sigma^2 = 50$, the process variance is \$4.5M(50) = \$225M

- Calendar year development
  - ◇ Rather than estimating the remaining IBNR for each accident year, we can estimate development for the next calendar year period beyond the latest diagonal
  - ◇ To estimate development for the next 12-month calendar period, we take the difference in growth functions at the two evaluation ages and multiply it by 1) the estimated ultimate losses for the loss development method OR 2) Premium*ELR for the Cape Cod method
  - ◇ The process variance and parameter variance are calculated as usual
  - ◇ A **major reason** for calculating the 12-month development is that the estimate is testable within a short timeframe. One year later, we can compare it to the actual development and see if it was within the forecast range

- Variability in discounted reserves
  - ◇ Use the same payout pattern and model parameters that were used with undiscounted reserves
  - ◇ The *CV* for discounted reserves is lower since the tail of the payout curve has the greatest parameter variance and also receives the deepest discount
  - ◇ See Appendix C section below for the calculation of discounted reserves, as well as an example

VI. Comments and Conclusion

- ◇ Abandon your triangles
  - The MLE model works best when using a tabular format of data (see exhibits in paper for an example) rather than a triangular format
  - All we need is a consistent aggregation of losses evaluated at more than one date
- ◇ The CV goes with the mean
  - If we selected a carried reserve other than the maximum likelihood estimate, can we still use the *CV* from the model?
    - ◇ Technically, the answer is "no". The estimate of the standard deviation in the MLE model is directly tied to the maximum likelihood estimate

- ◇ However, for practical purposes, the answer is “yes”. Since the final carried reserve is a selection based on a number of factors (some of which are not captured in the model), it stands to reason that the standard deviation should also be a selection. The output from the MLE model is a reasonable basis for that selection

◇ Other curve forms

- This paper focused on the loglogistic and weibull growth curves for a few reasons:
  - ◇ Smoothly move from 0% to 100%
  - ◇ Closely match the empirical data
  - ◇ First and second derivatives are calculable
- The method is not limited to these forms; other curves could be used

◇ **The main conclusion of the paper is that parameter variance is generally larger than the process variance**, implying that our need for more complete data (such as the exposure information in the Cape Cod method) outweighs the need for more sophisticated models

VII. Appendix B: Adjustments for Different Exposure Periods

◇ Before showing the final formula, let’s walk through a quick example:

- Assume we are 9 months into an accident year
- Then $G^*(4.5|\omega, \theta)$ represents the cumulative percent of ultimate of the 9-month period only (not the entire AY since a full AY exposure period is 12 months)
- In order to estimate the cumulative percent of ultimate for the full accident year, we must multiply by a scaling factor that represents the portion of the AY that has been earned
- Thus, the AY cumulative percent of ultimate as of 9 months is $G_{AY}(9|\omega, \theta) = (\frac{9}{12}) \cdot G^*(4.5|\omega, \theta)$

◇ Generalizing this process, there are two steps:

- **Step 1: Calculate the percent of the period that is exposed:**

  For accident years (AY):

$$Expos(t) = \begin{cases} t/12, & t \leq 12 \\ 1, & t > 12 \end{cases}$$

- **Step 2: Calculate the average accident date of the period that is earned:**

For accident years (AY):

$$AvgAge(t) = \begin{cases} t/2, & t \leq 12 \\ t-6, & t > 12 \end{cases}$$

◇ The final cumulative percent of ultimate curve, including annualization, is given by:

$$\boxed{G_{AY}(t|\omega,\theta) = Expos(t) \cdot G^*(AvgAge(t)|\omega,\theta)}$$

◇ *Note: Since the PY versions of the formulas above are unlikely to be tested, I have not included them*

VIII. Appendix C: Variance in Discounted Reserves

◇ Calculation of the discounted reserve, R_d:

$$\boxed{R_d = \sum_{AY}\sum_{k=1}^{y-x} ULT_{AY} \cdot v^{k-\frac{1}{2}} \cdot (G(x+k) - G(x+k-1))}$$

where $v = \frac{1}{1+i}$ and i is the constant discount rate

◇ Process variance of R_d:

$$\boxed{Var(R_d) = \sigma^2 \cdot \sum_{AY}\sum_{k=1}^{y-x} ULT_{AY} \cdot v^{2k-1} \cdot (G(x+k) - G(x+k-1))}$$

◇ **LDF method**

- For consistency, we will use the same LDF example shown earlier in the outline. Assume that expected loss emergence is described by a loglogistic curve. In addition, assume that the curve should be truncated at 120 months
- Given the following cumulative losses and parameters:

	Cumulative Losses ($)				
AY	12	24	36	48	60
2010	500	1500	2250	2590	2720
2011	550	1700	2400	2725	
2012	450	1200	2000		
2013	600	1750			
2014	575				

Parameters	
θ	21.4675
ω	1.477251
σ^2	59.9876

- We obtain the following results:

AY	Losses at 12/31/14	Age at 12/31/14	Avg. Age (x)	Growth Function	Fitted LDF	Trunc. LDF	Estimated Reserves	Estimated Ultimate
Trunc.		120	114	0.922	1.0846	1.0000		
2010	2720	60	54	0.796	1.2563	1.1583	430.576	3150.576
2011	2725	48	42	0.729	1.3717	1.2647	721.308	3446.308
2012	2000	36	30	0.621	1.6103	1.4847	969.400	2969.400
2013	1750	24	18	0.435	2.2989	2.1195	1959.125	3709.125
2014	575	12	6	0.132	7.5758	6.9848	3441.260	4016.260
Total							7521.669	17291.669

- Given a discount rate of 3%, let's determine the discounted reserves for AY 2011. To do this, we decompose AY 2011 into its CY pieces and discount them:

Age	Average Age	Growth Function	Trunc. LDF	Estimated Reserves	Discounted Reserves
Trunc.	114	0.922	1.0000	48.587	41.297
108	102	0.909	1.0143	59.892	52.433
96	90	0.893	1.0325	82.295	74.207
84	78	0.871	1.0586	115.676	107.436
72	66	0.840	1.0976	164.542	157.406
60	54	0.796	1.1583	250.315	246.643
48	42	0.729	1.2647		
				721.308	**679.421**

- Here are the calculations for age 72:
  - ◇ Avg. age $= 66 = 72 - 6$
  - ◇ Growth function $= \frac{x^\omega}{x^\omega + \theta^\omega} = \frac{66^{1.477251}}{66^{1.477251} + 21.4675^{1.477251}} = 0.840$
  - ◇ Trunc. LDF $= \frac{0.922}{0.840} = 1.0976$
  - ◇ Estimated reserves $= 3446.308\left(\frac{1}{1.0976} - \frac{1}{1.1583}\right) = 164.542$. This is the amount that emerges between ages 60 and 72

◇ Discounted reserves $= \frac{164.542}{1.03^{2-0.5}} = 157.406$. Since the average age is 66, the reserves must be discounted by 1.5 years to bring them back to the age 48 level

- *Please note that the sum of the estimated reserves over each CY piece (721.308) equals the estimated reserves found in the example shown earlier in the outline. This provides a nice check that we decomposed the reserves properly*

◇ **CC method**

- Given the following parameters for the CC method:

Parameters	
θ	22.3671
ω	1.441024
σ^2	50.0730

- As shown earlier in the outline, we obtain the following results:

AY	On-Level Premium	Age at 12/31/14	Average Age (x)	Growth Function	0.913 − Growth	Expected Losses	Estimated Reserves
Trunc.		120	114	0.913	0.000		
2010	5000	60	54	0.781	0.132	3490.00	460.680
2011	5200	48	42	0.713	0.200	3629.60	725.920
2012	5400	36	30	0.604	0.309	3769.20	1164.683
2013	5600	24	18	0.422	0.491	3908.80	1919.221
2014	5800	12	6	0.131	0.782	4048.40	3165.849
Total							7436.353

- Given a discount rate of 3%, let's determine the discounted reserves for AY 2011. To do this, we decompose AY 2011 into its CY pieces and discount them:

Age	Average Age	Growth Function	Estimated Reserves	Discounted Reserves
Trunc.	114	0.913	50.814	43.190
108	102	0.899	65.333	57.196
96	90	0.881	83.481	75.276
84	78	0.858	116.147	107.874
72	66	0.826	163.332	156.248
60	54	0.781	246.813	243.192
48	42	0.713		
			725.920	**682.976**

- Here are the calculations for age 72:
  - ◇ Avg. age = 66 = 72 − 6
  - ◇ Growth function = $\frac{x^\omega}{x^\omega+\theta^\omega} = \frac{66^{1.441024}}{66^{1.441024}+22.3671^{1.441024}} = 0.826$
  - ◇ Estimated reserves = 3629.6(0.826 − 0.781) = 163.332. This is the amount that emerges between ages 60 and 72. **Notice that we are multiplying the percentage to emerge by the expected losses, not the ultimate losses**. This is because the reserves for the CC method are based on the expected losses
  - ◇ Discounted reserves = $\frac{163.332}{1.03^{2-0.5}}$ = 156.248. Since the average age is 66, the reserves must be discounted by 1.5 years to bring them back to the age 48 level
- *Please note that the sum of the estimated reserves over each CY piece (725.920) equals the estimated reserves found in the example shown earlier in the outline. This provides a nice check that we decomposed the reserves properly*

