Does the Mack Method understate the range of possible outcomes?

Daniel Murphy, FCAS, MAAA
Trinostics LLC

CLRS 2009

GIRO WP Simulation Result

- Over 10% of the time, actual unpaid losses were greater than the Mack Method's 99th percentile
- "The main conclusion of the simulation work carried out so far is that **[the Mack Method tends] to understate the chance of extreme adverse outcomes**, even in situations where [its] underlying assumptions are perfectly satisfied."*
- "Commonly used [stochastic] methods are inadequate to cover the full range of reserving variability."
  – Morgan, K.A., et. al., "Actuarial Aspects of Internal Models for Solvency II"**

* GIRO WP Report, July 27, 2008, "Best Estimates and Reserving Uncertainty", http://www.actuaries.org.uk/__data/assets/pdf_file/0010/31303/BHPrize_Gibson.pdf

** Presented to the Institute of Actuaries February 23, 2009 http://www.actuaries.org.uk/__data/assets/pdf_file/0009/146664/sm20090223.pdf

GIRO WP Simulated 10,000 10x10 Triangles*

- Age 1 losses were independently sampled from a lognormal with mean = 1, var = 1
- Age 2 losses were independently sampled from a shifted lognormal with shift = X_1, mean = $X_1 b_1$ and var = $X_1\sigma^2$; similarly for age 3, …

b	4.2890	2.0640	1.5020	1.2680	1.1500	1.0850	1.0480	1.0270	1.0150	
σ^2	1	1	1	1	1	1	1	1	1	
AY \ age	1	2	3	4	5	6	7	8	9	10
2001	$0.420	$2.873	$7.175	$14.295	$16.676	$18.732	$36.983	$38.045	$38.102	$38.118
2002	0.824	3.742	8.875	12.354	14.176	14.680	15.044	15.080	15.111	15.347
2003	0.353	1.118	2.058	2.961	3.795	4.070	4.284	4.360	4.472	4.517
2004	2.669	8.403	12.937	18.463	24.133	24.811	25.665	25.725	25.729	25.873
2005	0.930	5.056	11.421	13.749	15.209	21.361	21.592	21.618	21.862	22.754
2006	0.357	1.382	2.485	3.002	3.135	6.299	6.455	6.523	6.550	6.565
2007	1.061	4.392	8.382	11.093	13.844	14.495	14.519	14.574	14.964	14.965
2008	1.308	6.626	10.563	14.934	20.307	20.861	26.047	26.345	26.398	26.406
2009	1.142	5.685	10.377	18.663	22.144	25.626	26.687	27.036	27.248	27.250
2010	1.639	7.667	16.534	27.154	37.583	40.970	64.295	64.772	65.366	65.411
									Sum	247.205
									Paid	136.731
									"Actual" Unpaid	**$110.474**

*Algorithm A: GIRO Report, pp. B1-B2

GIRO WP Algorithm A: Continued

- For each trial, the WP
  - Estimated the mean (expected value) unpaid loss amount via the chain ladder method using weighted average age-to-age factors (ata's)
  - Calculated the standard error (total risk) of the estimate using the Mack Method
  - Fit a lognormal distribution to the mean and standard error
  - Compared the 99th percentile of the fit to the previously simulated "actual" unpaid amount

$Y = Xb + \sqrt{X}\sigma z$

Trial 1
Estimated Unpaid Loss Amount

Accident Year	Current Diagonal	VW ata	LDF	Estimated Ultimate	Estimated Unpaid	*Simulated Actual Unpaid*
2001	$ 38.118		1.0000	$ 38.118	$ -	*$ -*
2002	15.111	1.0004	1.0004	15.118	0.006	*0.235*
2003	4.360	1.002	1.002	4.369	0.009	*0.157*
2004	25.665	1.021	1.023	26.255	0.590	*0.208*
2005	21.361	1.316	1.346	28.757	7.396	*1.392*
2006	3.135	1.131	1.522	4.772	1.637	*3.430*
2007	11.093	1.190	1.811	20.089	8.996	*3.872*
2008	10.563	1.423	2.578	27.227	16.664	*15.844*
2009	5.685	1.902	4.903	27.874	22.189	*21.565*
2010	1.639	4.334	21.249	34.827	33.188	*63.772*
Sum	$ 136.731			$ 227.406	**$ 90.675**	*$ 110.474*

Mack se		**39.085**	
	50%	90%	99%
Lognormal percentile	$83.269	$141.333	$217.550
Sufficient?	FALSE	TRUE	TRUE

$$Y = Xb + \sqrt{X}\sigma z$$

What happens over all 10,000 trials?

	Percentile			Simulated	Sufficient?		
Trial	50%	90%	99%	*Actual*	50%	90%	99%
1	$ 83.269	$ 141.333	$ 217.550	*$ 110.474*	FALSE	TRUE	TRUE
2	78.329	120.532	171.279	*140.618*	FALSE	FALSE	TRUE
3	59.184	76.377	94.028	*101.583*	FALSE	FALSE	FALSE
...	...	...	...	...	...	...	...
9999	134.400	177.211	222.018	*40.433*	TRUE	TRUE	TRUE
10000	47.780	79.419	120.181	*41.434*	TRUE	TRUE	TRUE

Percent sufficient	***59.0%***	***25.6%***	***10.2%***
WP Table B-3	58.30%	24.55%	10.1%

How can the Mack estimates be so far off?

1. Is the *mean* too low?
2. Is the *variance (MSE, standard error)* too low?
3. Is the lognormal too *thin-tailed*?
4. Something else?

Is the C.L. unpaid loss too low on average?

- Two-part question
- Part 1. How does the mean from the simulation compare with the expected value?

Accident Year	Expected Current Diagonal	Expected VW ata	Expected LDF	Expected Ultimate	Expected Unpaid	mean unpaid from simulations	% difference
2001	22.982		1.0000	$ 22.982	$ -		
2002	22.642	1.0150	1.0150	22.982	0.340	0.351	3.4%
2003	22.047	1.0270	1.0424	22.982	0.935	1.056	12.9%
2004	21.037	1.0480	1.0924	22.982	1.945	1.890	-2.8%
2005	19.389	1.0850	1.1853	22.982	3.593	3.567	-0.7%
2006	16.860	1.1500	1.3631	22.982	6.122	6.216	1.5%
2007	13.296	1.2680	1.7284	22.982	9.685	9.652	-0.3%
2008	8.852	1.5020	2.5961	22.982	14.129	14.155	0.2%
2009	4.289	2.0640	5.3583	22.982	18.693	18.839	0.8%
2010	1.000	4.2890	22.9816	22.982	21.982	22.284	1.4%
Sum	$152.394			$229.816	$77.422	$ 78.011	0.8%

- Not bad
- A better seed choice might improve AY 2003 accuracy

Is the C.L. unpaid loss too low on average?

- Part 2. How do the point estimates compare with the "actual" simulated unpaid losses, on average?

Accident Year	Expected Unpaid	mean actual unpaid (simulated)	% difference	mean predicted unpaid	% difference
2001	$ -	$ -		-	
2002	0.340	0.351	3.4%	0.427	21.6%
2003	0.935	1.056	12.9%	1.021	-3.3%
2004	1.945	1.890	-2.8%	2.134	12.9%
2005	3.593	3.567	-0.7%	3.598	0.9%
2006	6.122	6.216	1.5%	6.166	-0.8%
2007	9.685	9.652	-0.3%	9.610	-0.4%
2008	14.129	14.155	0.2%	14.590	3.1%
2009	18.693	18.839	0.8%	18.981	0.8%
2010	21.982	22.284	1.4%	22.185	-0.4%
Sum	$77.422	$ 78.011	0.8%	$ 78.711	0.9%

- Chain ladder point estimate looks reasonably close to the actual mean value
- Not too low on average

Is Mack total risk too low on average?

- Two-part question
- Part 1. How does the variance from the simulation compare with the expected variance?
  - Not capable of being answered at this time
  - Although there is a formula for the mean of Algorithm A projections, this author knows of no corresponding formula for the variance

$Y = Xb + \sqrt{X}\sigma z$

Is Mack total risk too low on average?

- Part 2. How does the variance of predicted unpaid loss compare with the actual variance?

Accident Year	std dev of actual unpaid	std dev of predicted unpaid	% difference
2001	$ -	-	
2002	2.637	8.589	225.8%
2003	6.200	9.149	47.6%
2004	6.954	23.875	243.3%
2005	8.198	9.761	19.1%
2006	12.553	12.643	0.7%
2007	14.163	13.348	-5.8%
2008	18.893	41.590	120.1%
2009	23.853	29.272	22.7%
2010	27.333	24.980	-8.6%
Sum	$ 46.757	$ 117.950	152.3%

- The spread of chain ladder projections appears to be more sensitive than the mean to random variation (surprising?)
- However, no firm evidence that the Mack method significantly understates the actual variability for "Algorithm A" type triangles

Are Algorithm A unpaid losses lognormally distributed?

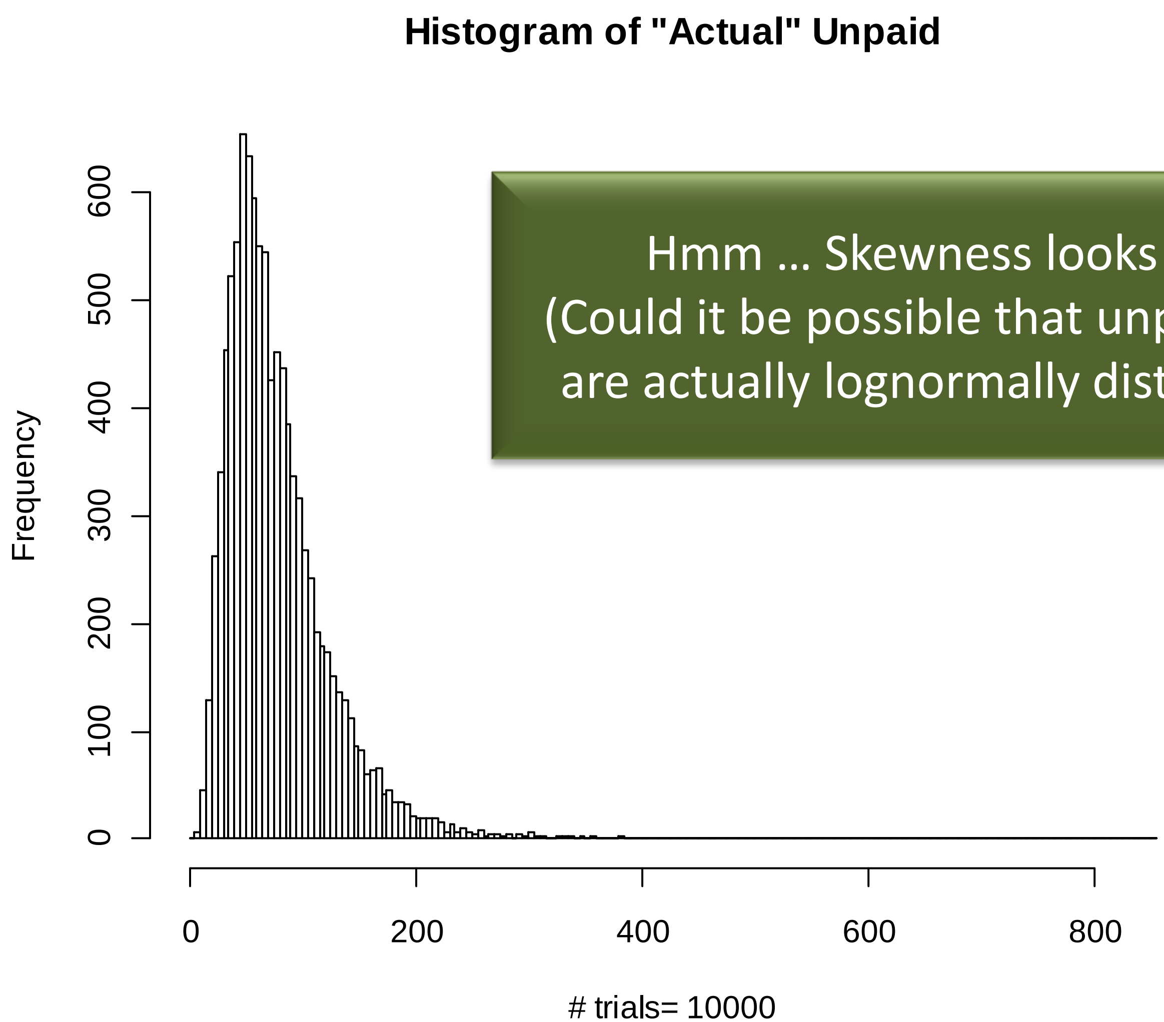

Hmm … Skewness looks right.
(Could it be possible that unpaid losses are actually lognormally distributed?)

Are Algorithm A unpaid losses lognormally distributed?

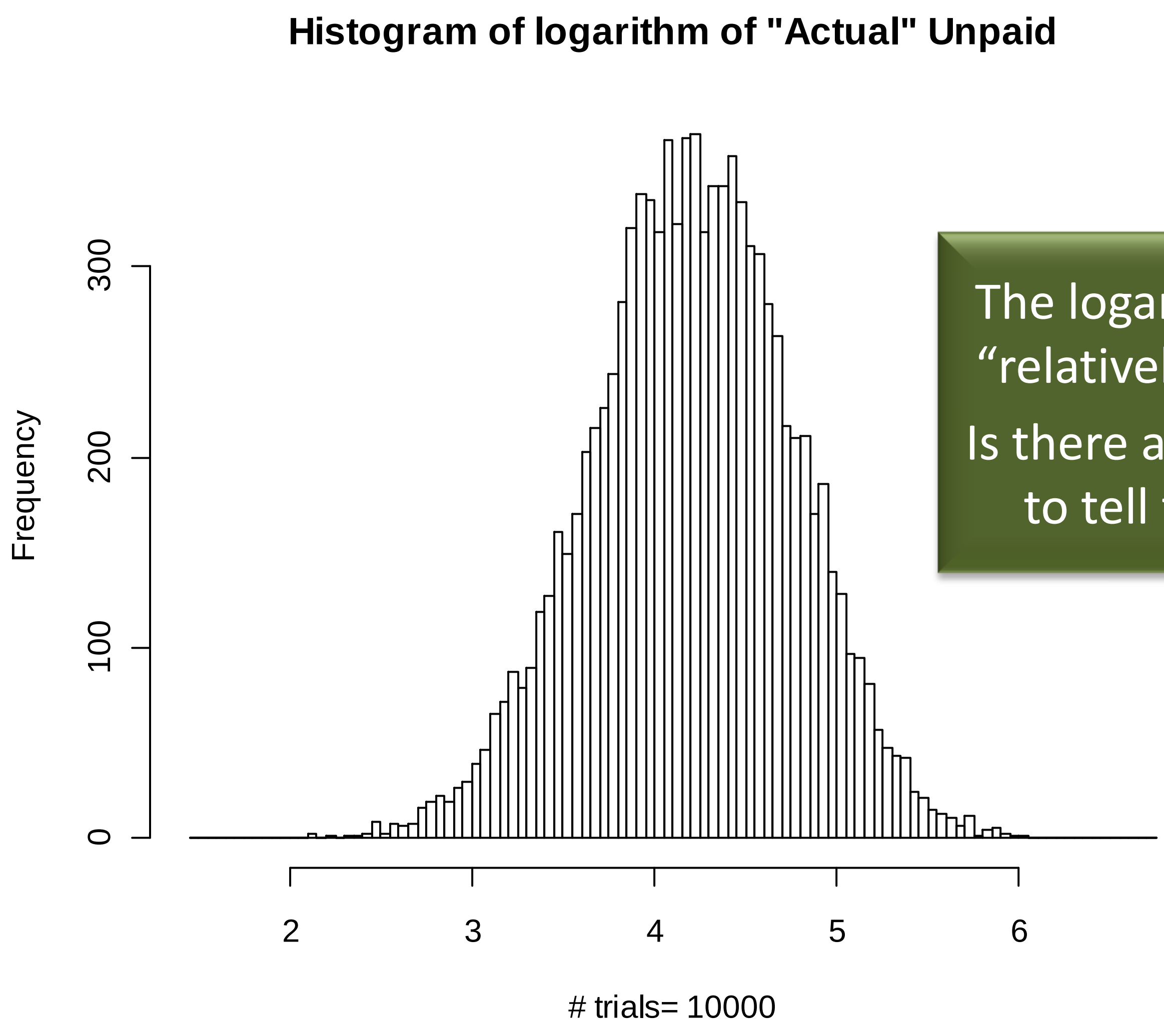

The logarithm looks "relatively" normal.

Is there a better way to tell for sure?

Are Algorithm A unpaid losses lognormally distributed?

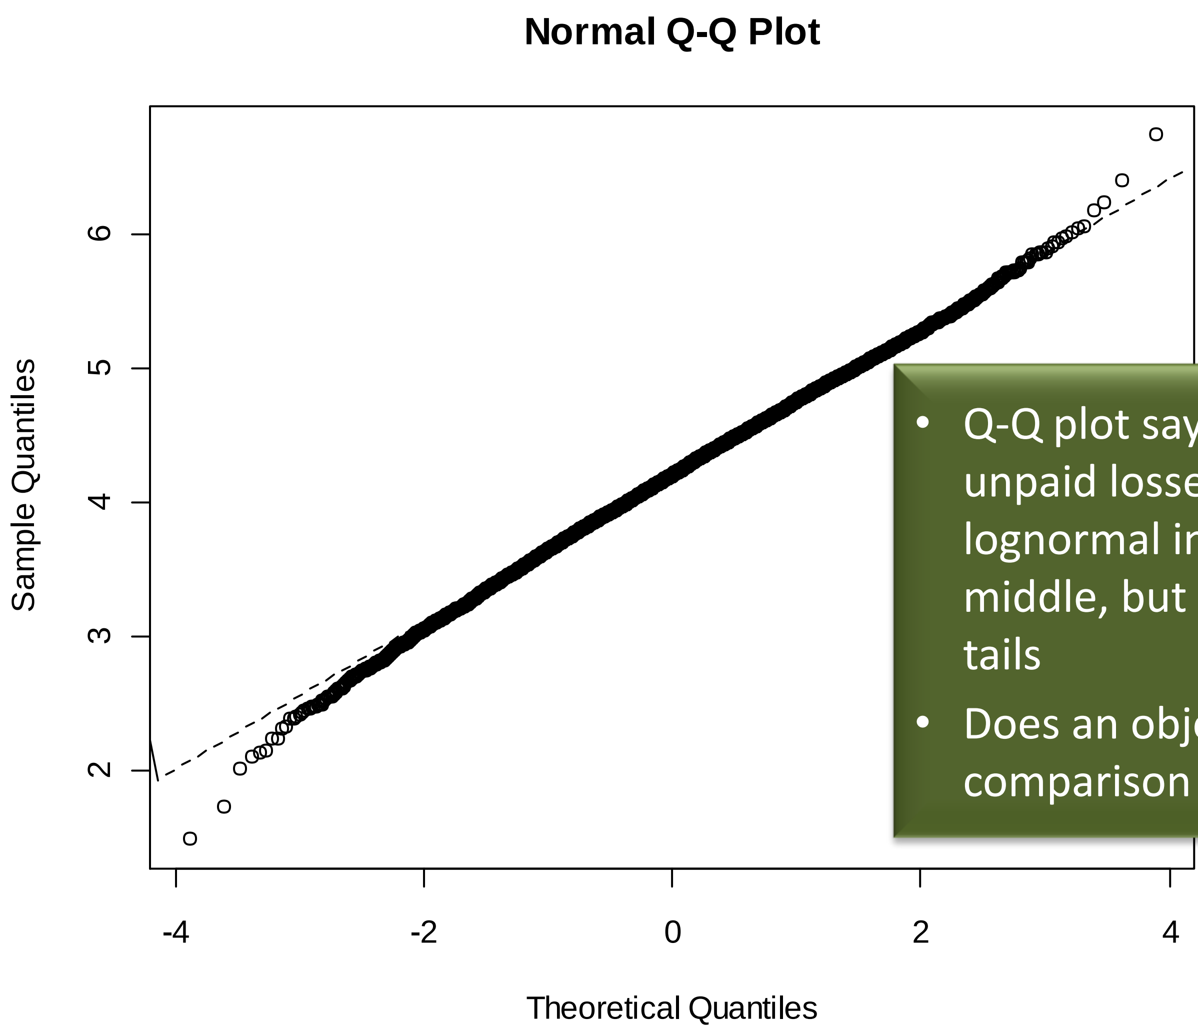

- Q-Q plot says actual unpaid losses look lognormal in the middle, but have fatter tails
- Does an objective comparison exist?

Are Algorithm A unpaid losses lognormally distributed?

- The Shapiro-Wilk test is a test of the hypothesis H_0 that the data come from a normal distribution
- In our situation, if the logarithm of unpaid losses passes the test, then the unpaid losses can be considered lognormally distributed
- For Algorithm A unpaid losses, the p-value = 0.001
- We reject H_0
  - (Interestingly, neither Kolmogorov-Sminov nor Wilcoxon tests rejected H_0)

Conclusion: Algorithm A unpaid losses do not follow a lognormal distribution

Taking stock

- The chain ladder method appears to produced reasonably unbiased point estimates for Algorithm A loss triangles
- The Mack Method standard error does not appear biased on the low side for Algorithm A loss triangles
- The lognormal distribution does not appear to be a good fit, particularly in the tails, for Algorithm A loss triangles
  – Recommended alternatives
    - Find a different distribution
    - Utilize simulation within the Chain Ladder Method
    - Avoid small triangles

Recommended alternatives

- Find a different distribution
  - The t distribution with Σ_{age}df degrees of freedom looks promising
- Utilize simulation within the Chain Ladder Method
  - Simulation is a practical alternative when the problem is too complicated for an analytical solution (cf. Gelman, *Data Analysis Using Regression and Multilevel/Hierarchical Models*, Cambridge, 2007)
- When WP's analysis was reproduced on triangles with 100 rows *with no other changes*, the percent of actuals that exceeded the estimated 99th percentile dropped from 10.2% to 2.7%

$Y = Xb + \sqrt{X}\sigma z$

Trinostics LLC is in the business of collaboration and education in the design and construction of valuable and transparent actuarial models.

Contact

Daniel Murphy, FCAS, MAAA
dmurphy@trinostics.com

  $Y = Xb + \sqrt{X}\sigma z$