Correlation and the Aggregation of Unpaid Loss Distributions

Paul J. Brehm, FCAS, MAAA

CORRELATION AND THE AGGREGATION OF UNPAID LOSS DISTRIBUTIONS

PAUL BREHM, FCAS, MAAA

Abstract

Significant progress has been made in the last decade in developing models to describe the distribution of unpaid losses for a line of insurance. Line-by-line distributions must be aggregated, however, in order to address company-wide issues such as enterprise risk, capital requirements, fair value, and more.

Using U.S. industry commercial lines data, this paper uses Zehnwirth's method to produce distributions of unpaid losses by line of insurance and Wang's standard normal copula method of aggregating correlated risk portfolios to create aggregate distributions of unpaid losses. In doing so, a methodology for direct estimation of correlations between lines is proposed.

1. MODEL OF UNPAID LOSSES[1]

Means. Actuarial loss reserving training is focused on getting a decent estimate of the mean of the unpaid loss and loss expense liability. Witness the CAS syllabus does not contain one reading concerning models for unpaid losses that reflect the stochastic nature of unpaid losses, producing distributions, or confidence intervals. Perhaps that's fair: it is certainly necessary to get the mean right. It's necessary, but not sufficient.

ASOP 36 speaks of "materiality," "ranges of reasonable estimates," "adverse deviations," and "risk margins." Statutory codification refers to "best estimates" and "risk based capital." The IASB speaks of "fair value liabilities." All of these concepts have a common element: they require a view of the higher moments of an unpaid loss distribution. We are, after all, in the business of risk, and to truly understand the risk associated with a portfolio of loss reserves the actuary must have a view about the distribution of unpaid losses, not just the mean.

Higher moments. A number of methodologies have been proposed to produce confidence intervals about a mean reserve estimate, or alternatively, produce an estimated distribution of unpaid losses. In particular the 1994 call paper program on the variability of loss reserves and subsequent papers published some truly landmark methodologies: Halliwell [9], Holmberg [10], Mack [12], Murphy [15], Verall [20], and Zehnwirth [24] just to name a few. While each model has its strengths and weaknesses, this paper

[1] Throughout this paper the term "loss" is used to refer to all loss and defense costs and adjusters fees (the old allocated loss adjustment expenses) for simplicity.

applies Zehnwirth's model to industry Schedule P data to illustrate the estimation of a distribution of unpaid losses.

Zehnwirth example. The data used in the following example was taken from the U.S. total industry 2000 Schedule P as compiled by A.M. Best[2]. Commercial lines, excluding excess of loss reinsurance lines, were analyzed. Data was grouped into six lines (commercial auto, workers' compensation, commercial multi-peril, medical malpractice, all other liability, and all other[3]) in an attempt to minimize detail.

The industry (non-medical) liability triangle will be used to illustrate Zehnwirth's modeling framework. Graph 1.1 shows the cumulative development by accident year for these data.

Graph 1.1

Industry Liability Cumulative Paid Losses

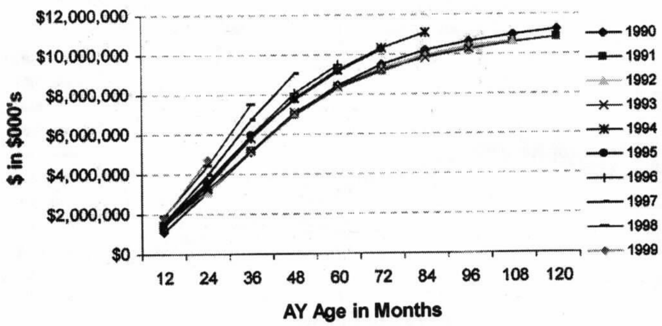

Zehnwirth's model can be characterized by the following algorithm.

1. Start with a cumulative paid loss data triangle (illustrated above).
2. Calculate an incremental paid triangle.
3. Inflation adjust the incremental paids, if desired.
4. Preferably divide each incremental paid by exposure, or in the absence of exposure as with industry data, by earned premium.
5. Calculate the natural logs of the incremental ratios in step #4.

---

[2] The 1990 accident year was included, too, at least through year-end 1999, from the 1999 report. Current reserves for the 1999 accident year were estimated by running off a portion of those carried at year-end 1999.

[3] The 'all-other' -- principally short tailed -- triangle was created by subtracting each reported line triangle (including homeowners and private passenger auto liability) from the total, all-lines triangle. This does not yield a pure short tailed commercial lines triangle. The triangle will still include personal auto physical damage and inland marine floaters. This triangle was included anyway, for illustrative purposes.

6. Model the log incremental paid ratios with a regression model, incorporating dummy variables for the accident year, development year, and calendar year dimensions of the triangle.

Focus is on paid data because logs are used on the incremental ratios. Since the increments cannot be negative, downward development is problematic. In fact, even paid data can present a problem in lines where negative incremental paids are common (e.g., surety).

Graph 1.2 shows the transformed triangle for the cumulative data in Graph 1.1: the log of incremental paid losses to net earned premium. The data was not inflation adjusted, but could be without loss of generality.

Graph 1.2

Log Incremental Paid/NEP

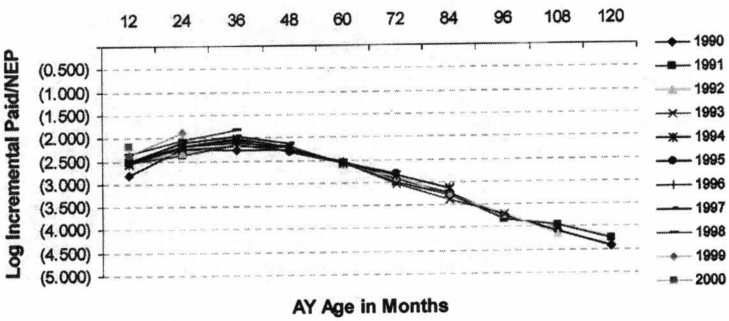

Zehnwirth's model describes the above data patterns with parameters for the accident year ("i") dimension (α_i, which are essentially the vertical leveling of the accident year), the development year ("j") dimension (γ_j, which actuaries would call the incremental payout pattern), and the calendar year (i+j) dimension (ι_{ij}, which would represent some sort of calendar period distortion not otherwise picked up by the parameters in the other two dimensions, e.g., a shift in inflation or a change in the claims department). The parameterization is such that tail development is already calculated as part of the assumed exponential decay by the last fitted γ. The general formula, then, looks as follows. For the log incremental paid ratio for accident year i and development period j:

$$y_{ij} = \alpha_i + \Sigma^{j}_{k=1}\gamma_k + \Sigma^{i+j}_{t=1}\iota_t + \varepsilon_{ij} \qquad 1.1$$

In essence, Zehnwirth's model describes the log of each incremental payment ratio by the combination of three vectors, one each for the accident year, the development period, and the calendar year.

The regression model (1.1) assumes that the ε will be normally distributed. Thus, each estimated incremental payment, when transformed back into dollars, will be distributed logormally. Since the distribution of the sum of lognormals is not a simple closed form distribution, the aggregate unpaid loss must be simulated in a conditional simulation using formula 1.1 above.

Graph 1.3 shows the simulated CDF for the industry liability unpaid losses. A lognormal curve was fit to the simulated distribution using a simple method-of-moments fit. Despite the fact that the sum of lognormals theoretically isn't lognormal, a lognormal curve fits the data well.

Graph 1.3
Industry Commercial Liability Unpaid Loss CDF

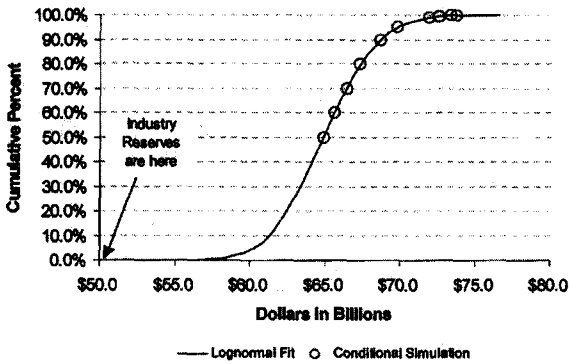

Loss triangles for each commercial line were modeled as above. Distributions were estimated using conditional simulations with formula 1.1. As with the liability example, lognormal models were fit to each distribution with good results. Table 1.1 below summarizes the results.

Table 1.1

U.S. Industry Commercial Lines Reserves 1990-2000

	Model Unpaid	Industry Carried	Redundancy /(Deficiency)	% tile	Lognormal Parameters μ	σ
Commercial Auto [C]	23,008	18,911	(4,097)	0.0%	3.135	0.032
Work Comp. [D]	66,535	60,597	(5,938)	15.8%	4.194	0.089
CMP [E]	27,734	24,753	(2,981)	0.0%	3.322	0.018
Med Mal [F]	26,098	19,478	(6,620)	0.0%	3.261	0.032
Liability [H,R]	64,950	50,148	(14,802)	0.0%	4.173	0.045
All Other	26,254	24,578	(1,676)	27.0%	3.263	0.099
	234,579	198,465	(36,114)	?	?	?

Table 1.1 highlights the central issue addressed in this paper. While estimating the distributions of unpaid losses by line is an important advance in actuarial science, the actuary is still ill equipped to estimate the distribution of unpaid losses for the entire portfolio.

We could apply the Zehnwirth model to an industry total, all lines triangle and get both a mean result and an associated distribution. As actuaries, we bristle at the thought: the lack of homogeneity in such a triangle would surely distort our ability to get the mean unpaid loss estimate correct. So, we model unpaid losses in more homogenous, hopefully still credible, segments. Line-by-line we can calculate the mean unpaid loss and its distribution, but now we have lost the ability to calculate a simple closed form distribution for the aggregate unpaid loss reserve requirement. We cannot fill in the empty boxes in Table 1.1 without a methodology for aggregating correlated distributions.

2. AGGREGATE DISTRIBUTIONS

Correlations. Methods exist for aggregating distributions, but the correlations between the distributions are always the critical component. Actuaries seem to appreciate the potentially profound impact correlations have on required economic capital, capital allocations or risk loads, reinsurance buying, etc., but very little is written about it[4].

Polar cases of correlations are simple and illustrative. For example, given marginal distributions of unpaid losses for two lines of business, an aggregate distribution can be easily created if one assumes the correlation is –1, 0, or 1 between the lines. If correlation is 0, one could simulate an amount from each of the distributions and simply add the two numbers together. If the correlation is 1, amounts are simulated for each distribution, sorted, and matched up from smallest to largest, and then added together (for –1 match opposite rank orders). Alternatively, an aggregate distribution can be computed in closed form with a simple variance-covariance matrix and an assumption regarding functional form.

As a basis, the marginal lognormal distributions shown in Table 1.1 were aggregated in just such manners for the polar cases of no correlation and perfect correlation. The results are shown in graph 2.1 below.

Graph 2.1
Aggregate Unpaid Loss Distributions
with Polar Correlation Assumptions

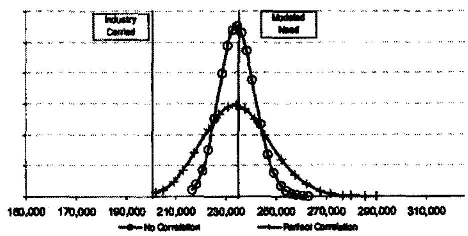

---

[4] Wang [22] mentions a constant correlating factor for unpaid losses in an appendix. Myers [23] discussion of Wang illustrates the incorporation of correlation into the calculation of aggregate loss distributions in a collective risk theory framework. The correlation parameter behaves much like the contagion parameter of the collective risk model. Frequency is the correlating factor between lines.

In conversations with Todd Bault, he has hypothesized a model of correlation wherein correlation behaves as a scaling factor. His argument goes something like this: if you believe that everything is correlated to the whole, then the larger the line, the higher the correlation. Chris Gross has offered a similar model, one where there exists a correlation within a line of subsets of that line (e.g., a new account). While currently unpublished, I think these views have some merit and should be explored.

The public access DFA model [5] addresses correlations in unpaid losses by essentially injecting them into future loss payments. Simple linear models are used to describe line of business loss cost inflation as a function of the CPI. Thus all lines are correlated with each other since they are based on the CPI. This correlating inflation is applied in the accident year dimension.

Graph 2.1 clearly shows the importance of correlation assumptions in computing an aggregate distribution – not for the mean – but for the risk component. If the industry were required to hold capital to, say, the 99.9th percentile over carried reserves, required economic capital would be roughly $60 billion for the uncorrelated case and over $80 billion for the perfectly correlated case.

Correlations are likely not zero or one. Nor are we typically dealing with only two lines. For correlations other than the polar values, or for more than two lines, mixing distributions is no longer so simple.

Measuring correlations. The single biggest source of risk in an unpaid loss portfolio is arguably the potential distortions that can affect all open accident years, i.e., changes in calendar year trends. Such a distortion could be a social inflation, say in court judgments, that affect all open claims (subsequent incremental payments) at once. In Zehnwirth's model, this is saying that the future ι_{ij}'s turn out to be much different than predicted and reserved for.

Our line-by-line distributions reflect the calendar year inflation risk, since ι_{ij} is a statistical estimate and has an associated variability that is incorporated into the conditional simulations. But what happens if a calendar year distortion affects more than one triangle simultaneously? This is the sort of thing that keeps chief actuaries awake at nights: the fear that a systematic distortion will affect multiple lines, in the same (bad) direction all at once.

Having modeled a number of lines of business, we have a vector of calendar period trend parameters, ι_{ij}, for each line. By measuring the line-by-line correlations between these historic parameters, we can estimate a correlation matrix for calendar period movements. The estimated calendar year parameters and associated correlation matrix from this exercise are shown below[5].

[5] It's worth mentioning correlation measures that don't work for the purposes of unpaid loss distributions. Correlations in future unpaid losses cannot be calculated by looking at time series of loss ratios. Historic loss ratios are typically highly correlated because of cyclical pricing impacts in the denominator of the loss ratio. For loss reserving purposes pricing induced correlations do not matter. Furthermore, correlations in the numerator of a loss ratio, if measured at an ultimate value are suspect depending on methodology. If, for example, two lines are developed to ultimate using the Bornhuetter-Ferguson methodology with the same or similar seed loss ratios, correlation is actually injected into the "data."

Figure 2.1

Calendar Period Inflation Parameter Estimates (ι_{ij})

	Commercial Auto [C]	Work Comp. [D]	CMP [E]	Med Mal [F]	Liability [H, R]	Other
1991-1992	9.18%	3.70%	-5.96%	0.00%	0.00%	0.00%
1992-1993	5.09%	3.70%	2.67%	0.00%	0.00%	0.00%
1993-1994	5.09%	3.70%	2.67%	0.00%	0.00%	0.00%
1994-1995	5.09%	3.70%	-3.29%	0.00%	0.00%	0.00%
1995-1996	5.09%	3.70%	-3.29%	0.00%	0.00%	0.00%
1996-1997	2.70%	3.70%	-3.29%	-5.64%	0.00%	0.00%
1997-1998	2.70%	10.24%	-3.29%	6.06%	-3.25%	0.00%
1998-1999	5.74%	0.00%	-3.29%	6.06%	5.34%	0.00%
1999-2000	5.74%	15.05%	-3.29%	6.06%	5.34%	5.01%

Correlation in the historical calendar year inflation trends is evident by inspection of the above table. While not uniformly true, it is apparent that inflation is low and remarkably stable through out the middle-90's. Inflation even inexplicably declines in some lines in or near 1996. Finally, there seems to be evidence that inflation accelerates some time in 1998 or 1999.

To compute the correlation matrix from the above table, simply calculate all pair-wise correlations between the lines of business. In Excel, for the correlation between auto and work comp for example, this would be the formula: =correl(AUTO, COMP). These calculations were made and are shown below in Figure 2.2.

Figure 2.2

Correlation Matrix

	Commercial Auto [C]	Work Comp. [D]	CMP [E]	Med Mal [F]	Liability [H,R]	Other	average
Commercial Auto [C]	1.000	(0.169)	(0.259)	0.100	0.337	0.115	0.025
Work Comp. [D]	(0.169)	1.000	(0.138)	0.465	0.079	0.812	0.210
CMP [E]	(0.259)	(0.138)	1.000	(0.139)	(0.118)	(0.132)	(0.157)
Med Mal [F]	0.100	0.465	(0.139)	1.000	0.396	0.444	0.253
Liability [H,R]	0.337	0.079	(0.118)	0.396	1.000	0.611	0.261
Other	0.115	0.812	(0.132)	0.444	0.611	1.000	0.370
							0.160

The average correlation across all values is 0.16 (ignoring the diagonal of the correlation matrix). While fairly small, it is obviously the result of offsetting negative and positive values. The above matrix was accepted as-is in further examples below, but could be judgmentally amended to cap correlations at a certain level, set small (statistically insignificant) correlations to zero, or to erase correlations that lack an intuitive reason for being. No consideration was given to statistical significance, but this may be a fruitful area for future research.

This methodology provides a simple and practical method to measure correlations of unpaid losses between lines of business given typical actuarial data arrays. It is dependent on a loss reserving model capable of estimating calendar year trends, in this case the Zehnwirth model.

The method proposed above is not without its shortcomings.

1. Such a framework is that is inherently a slave to the data at hand. The decade of the 90's, used here for illustration, was a decade marked by low, stable inflation in a prosperous economy. Correlations measured in such an environment will reflect that environment and perhaps nothing more. Actuarial judgment should play a role in adjusting assumptions in changing or different environments.

2. The calendar year inflation parameters are themselves estimates from a model. There is often no one 'true' model, and different model parameters will yield different correlation results.

3. This model measures correlations in unpaid loss distributions by asserting that the principal correlating factor is severity. There is no consideration of frequency, i.e., correlations between lines in pure IBNR claims.

There are undoubtedly additional critiques that could be raised. But, as my father used to say, "it's way better than what comes in second place" – which, from a review of available literature, appears to be nothing.

Between line correlations could be estimated directly if we could model all lines simultaneously with regression model 1.1. By estimating parameters for all lines at once, the variance-covariance matrix from the regression would reflect the correlations between lines. But models and computing power aren't quite there yet. In the mean time the method proposed here serves as a proxy.

Aggregation of Line Distributions – Model 1. Given distributions by line and a correlation matrix, an aggregate distribution can be created numerically using a standard normal copula[6]. Shaun Wang [21] proposed this method in his paper, "Aggregation of Correlated Risk Portfolios: Models and Algorithms," PCAS LXXXV, 1998 (pp. 887-891)[7].

[6] The standard normal copula has been criticized recently for producing uncorrelated tail values, which would clearly defeat the purpose at hand. Cf, Mango [14].
[7] Wang's method is not the only such method. Nakada [16] uses a numerical integration routine that is mathematically similar and does not require simulation.

Wang's standard normal copula algorithm is as follows.

1. Measure the correlation matrix Σ with elements ρ_{ij} (above).

2. From Σ construct the lower triangular matrix **B** via the Cholesky decomposition such that $\Sigma = \mathbf{BB'}$. Each element of B can be defined by:

$$b_{ij} = \frac{\rho_{ij} - \sum_{s=1}^{j-1} b_{is} b_{js}}{\sqrt{1 - \sum_{s=1}^{j-1} b_{js}^2}} \qquad 1 \le j \le i \le n \ \& \ \Sigma^0 = 0$$

3. Generate a vector $\mathbf{Y} = (y_1, \ldots, y_k)'$ of standard normals , where k=# of lines,. This is just the Excel function NORMINV(RAND()).

4. Define $\mathbf{Z} = \mathbf{BY}$. $\mathbf{Z} = (z_1, \ldots, z_k)'$ has the appropriate joint pdf defined by the correlation matrix. In Excel this is the array function MMULT(**B**,**Y**).

5. Set $u_i = \Phi(z_i)$, where Φ denotes the standard normal cdf. In Excel, NORMDIST(z_i,0,1,TRUE).

6. Set $x_i = F_i^{-1}(u_i)$, where F_i is the marginal distribution function for the modeling line i. In our case the F_i's are the assumed lognormal distributions produced by ICRFS. However, there is no restriction on the marginals. In fact, F need not be parameterized. An empirical distribution can be used. If a lognormal is used, the Excel formula is LOGINV(u_i,μ_i,σ_i)

7. Iterate steps 3-6 as many times as desired.

Wang shows that the standard normal copula methodology has the nice properties of creating a distribution of values with the desired correlations and still retaining the original marginal distributions. Furthermore, the required calculations are easily accomplished in a spreadsheet. If there is a drawback to this numerical methodology, it lies in the requirement to simulate to calculate the distribution. Inherent to all such methodologies, the subsequent calculations will not typically replicate the original estimated distribution.

For the industry example, the empirical pdf, based on 3,000 simulations, of the aggregate unpaid loss distribution looks as follows.

Graph 2.2

Distribution of Unpaid Losses

U.S. Commercial Lines Total

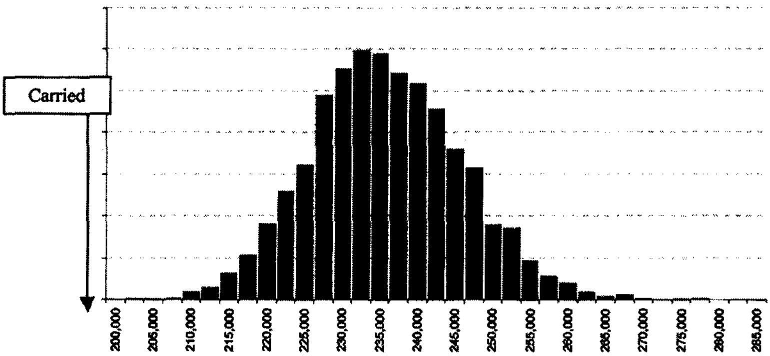

The minimum value in the above distribution is \$200 billion. The maximum value is \$276 billion. (Recall that the mean is \$234.5 billion.) The observed aggregate standard deviation is \$10 billion. The industry carried reserves for commercial lines loss and allocated loss expense from 1990 to 2000, at \$198.4 billion, are below the scale based on the above aggregate distribution.

Aggregation of Line Distributions – Model 2. Model 2 is a quick-and-dirty alternative. The mean of the aggregate distribution is known. It is simply the sum of the line means. The variance of the aggregate distribution can be calculated from the estimated variance-covariance matrix (VCM). The aggregate variance is the sum of the elements in the matrix. Given the aggregate mean and variance, a distribution can be estimated by assuming an appropriate functional form, e.g., lognormal. Following are the computations for the running example to compute a method of moments lognormal from the observed data.

Given the standard deviations calculated by line, create a matrix, σ =diagonal (σ_1, σ_2, …,σ_6). With the correlation matrix , Σ, the VCM will take the form:

$$\text{VCM} = \sigma'\Sigma\sigma$$

The VCM from the running example is this paper is shown below.

Figure 2.3

Variance-Covariance Matrix

	Auto [C]	Comp. [D]	CMP [E]	Med Mal [F]	Liability [H,R]	Other
Auto [C]		(0.75)	(0.10)	0.06	0.73	0.22
Comp. [D]	(0.75)		(0.41)	2.31	1.36	12.66
CMP [E]	(0.10)	(0.41)		(0.06)	(0.17)	(0.17)
Med Mal [F]	0.06	2.31	(0.06)		0.96	0.97
Liability [H,R]	0.73	1.36	(0.17)	0.96		4.63
Other	0.22	12.66	(0.17)	0.97	4.63	

The aggregate mean, M, is $234.5 billion. The aggregate standard deviation, S, is $9.8 billion. The method of moments estimates for lognormal parameters μ and σ of a distribution defined by observed moments M and S are 12.364 and 0.043 respectively. This lognormal distribution from method 2 is graphed below along with the results from the copula and the lognormal fit to the copula results.

Graph 2.3

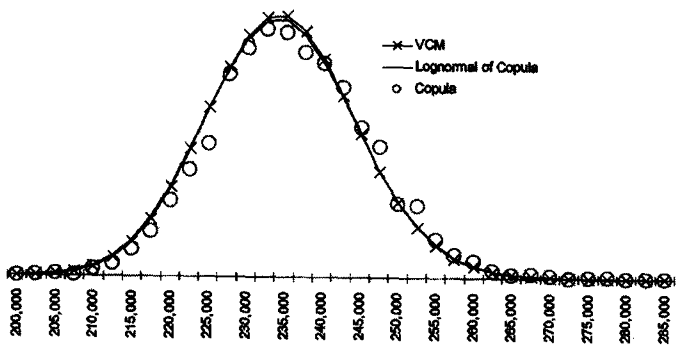

3. FUN WITH DISTRIBUTIONS

Given distributions of unpaid losses by line and in total, there are a number of interesting uses and implications. This section covers some topics where distributions are valuable in actuarial practice.

Capital requirements. Economic capital requirements line-by-line could be established using a ruin theory or value-at-risk construct. How much capital is needed to be sure that there are enough funds for claimants at some extreme probability? For example, at a 3 in 10,000 ruin probability (the equivalent of a AA credit rating default value) the required risk based capital, by line is the difference between the carried reserve and the 99.97th percentile (F^{-1}(0.9997)-reserve):

Table 3.1

Capital Requirements

	Carried Provision	$F^{-1}(0.9997)$	Stand Alone Capital	Reserves-to-Capital
Commercial Auto [C]	18,911	25,694	6,783	2.79
Work Comp. [D]	60,597	90,057	29,460	2.06
CMP [E]	24,753	29,496	4,743	5.22
Med Mal Claims made [F2]	19,478	29,107	9,629	2.02
General Liability Occurrence [H1]	50,148	75,605	25,457	1.97
Short Tailed Lines	24,578	36,745	12,167	2.02
Sum Total	198,465	286,704	88,239	2.25
Modeled Aggregate	198,465	271,161	72,696	2.73

The stand alone and aggregate capital needs shown above include a provision to cover the reserve inadequacies. If booked reserves were increased, supporting capital would decline, and the reserves-to-capital ratios shown above would increase.

The sum of the stand-alone capital amounts shown above, $88 billion, is a meaningless number: it is more economic capital than is required. To calculate the required economic capital to support the industry reserves, the total industry distribution must be used. This distribution reflected the portfolio aggregation and diversification effects for the industry. Required aggregate economic capital is shown in the table above. Its derivation is shown graphically below.

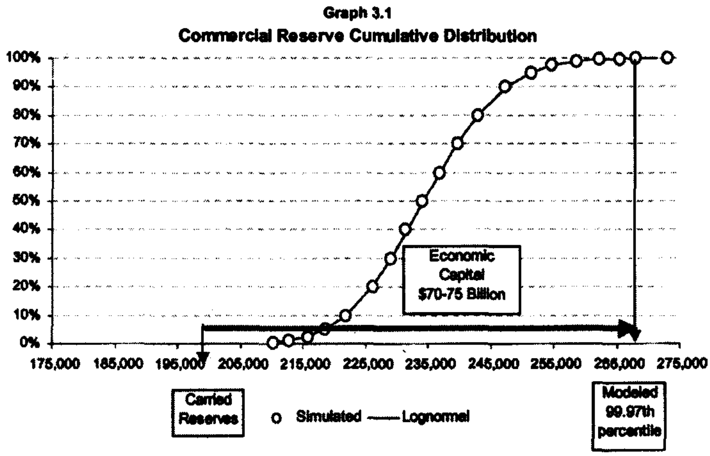

Graph 3.1 implies a required economic capital of roughly $73 billion. The sum of the line-by-line, stand-alone capital was $88 billion. Thus, there is a $15 billion natural diversification effect[8].

This example highlights a flaw in standard rules of thumb, such as reserves-to-surplus or premium-to-surplus. Though widely accepted, they mismatch numerator (partial economic risk factor) and denominator (accounting based total surplus).

Capital allocation. It was not the intent of this paper to dive into the capital allocation debate. However, given an aggregate risk distribution as described in the previous section, and given the variance-covariance matrix that necessarily underlies it, the source information exists to allocate capital to an unpaid loss portfolio in a defensible fashion.

A few years ago there was a terrific discussion thread in the CAS Proceedings, starting with Shalom Feldblum [7] and ending with Todd Bault [8] thrashing about whether one should allocate capital (risk load) in proportion to standard deviation or variance. Bault proved that the two (and actually others as well) were part of a broader, unified theory of risk based on correlations. When correlation approaches zero, the allocation is in

[8] This capital requirement is not the same as the total capital required for the commercial lines industry. The analysis would have to incorporate additional risks: volatility in new business (UPR), cats, investment risk, operational risks, etc. This could be easily accomplished. We need only to measure risk distributions for each risk type and use the same integration routine to get a total aggregate risk distribution (see [17]).

proportion to variance. When correlation approaches one, the allocation basis tends toward relative standard deviation.

Bault's conclusions immediately present a plausible capital allocation mechanism, incorporating standard deviation/variance and all of the measured covariances.
Zero and one are obviously just the polar cases. The more general result can be seen with a simple two-by-two VCM for risks 1 and 2:

$$\begin{pmatrix} \sigma_1^2 & \rho_{12}\sigma_1\sigma_2 \\ \rho_{12}\sigma_1\sigma_2 & \sigma_2^2 \end{pmatrix}$$

Given the above variance-covariance matrix, the capital allocation to a line is simply the sum of a given row in the variance-covariance matrix divided by the sum of the entire matrix. This generalized construct accounts for those cases where lines have perfect correlation, no correlation, or anything in between. And it is computationally tractable with any number of lines.

In the simplified VCM above for risk 1 and risk 2, the allocation formula for, say risk 1, would be

$$\frac{\sigma_1^2 + \rho_{12}\sigma_1\sigma_2}{\sigma_1^2 + \sigma_2^2 + 2\rho_{12}\sigma_1\sigma_2}$$

It is easy to see from this formula that, if ρ equals 0, the allocation basis is variance, and, if ρ equals 1, the allocation basis is standard deviation.

In the industry example, the resulting capital allocation looks as follows.

Figure 3.2
Capital Allocation from the VCM

	Commercial Auto [C]	Work Comp. [D]	CMP [E]	Med Mal [F]	Liability [H, R]	Other	Row Sum	Capital Allocation %
Commercial Auto [C]	0.6	(0.7)	(0.1)	0.1	0.7	0.2	0.7	0.7%
Work Comp. [D]	(0.7)	35.5	(0.4)	2.3	1.4	12.7	50.7	52.4%
CMP [E]	(0.1)	(0.4)	0.2	(0.1)	(0.2)	(0.2)	(0.7)	-0.7%
Med Mal [F]	0.1	2.3	(0.1)	0.7	1.0	1.0	4.9	5.1%
Liability [H,R]	0.7	1.4	(0.2)	1.0	8.4	4.6	15.9	16.4%
Other	0.2	12.7	(0.2)	1.0	4.6	6.8	25.1	26.0%
							96.7	100.0%

Our numerical example yields a negative capital allocation, which, while problematic, is not unique. Myers and Read [16] suffer the same problem. For the purposes of this paper, the observation is noted and accepted, in the spirit of presenting methodology first, an illustration second, and specific parameters a distant third.

The allocated capital and associated reserves-to-capital ratio are shown below, along side the stand alone capital estimates. In this exhibit, the total required capital and the allocated amounts are predicated on reserves being booked at an adequate level according to the analysis presented in this paper.

Figure 3.3

Capital Allocations for Unpaid Losses

	Required Reserves	Capital Allocation %	Allocated Capital	Reserves-to-Capital	Stand Alone Capital	Reserves-to-Capital
Auto [C]	$23,008	0.7%	$286	80.55	$2,686	8.57
Comp. [D]	$66,535	52.4%	$20,146	3.30	$23,522	2.83
CMP [E]	$27,734	-0.7%	-$261	(106.22)	$1,762	15.74
Med Mal [F]	$26,098	5.1%	$1,962	13.30	$3,009	8.67
Liability [H, R]	$64,950	16.4%	$6,314	10.29	$10,655	6.10
Short	$26,254	26.0%	$9,999	2.63	$10,491	2.50
	$234,579	100.0%	$38,446	6.10	$52,125	4.50

Valuation methodology. When placing an intrinsic or actuarial value on a prospective acquisition target, actuaries look at pricing 1) the true economic value imbedded in the balance sheet, and 2) the present value of the ongoing business. The method essentially turns everything past, present, and future into tangible (but stochastic) cash flows, capitalizes them appropriately, and then calculates a present value[9].

In light of the diversification benefit observed in the industry example above, it is apparent that valuation methodologies cannot attempt to measure the 'true value' of a target company in a vacuum, i.e., on a stand-alone basis. There exists a portfolio effect that will allow the merging companies to free up some amounts of capital for other investments or return to shareholders.

This free capital flow should be accounted for in the valuation, presuming that the theoretical diversification effect can be harvested given regulatory restrictions. Or should it? There is an open and active debate on this very subject in actuarial circles today. The opposing view would say that mergers between companies cannot create capital efficiencies that the market would not have already anticipated or that market instruments could not replicate.

[9] As an alternative, one can discount the cash flows at a risk-adjusted rate. As another, perhaps best, alternative, one can compute the risk adjusted distribution of cash and then discount at the risk free rate.

Fair Value Liabilities. The IASB has proposed that all liabilities, like assets, should be marked to market for financial reporting purposes. Lacking an active market of unpaid loss liabilities, a "fair value" must be estimated based on cash flows and risk. In the absence of an aggregate distribution as presented here, the total reserve portfolio risk cannot be quantified. Further, it should be apparent that it is insufficient to conduct a fair-value-ing exercise on a line-by-line basis, as it misses the diversification effect.

In the absence of a liquid market, think of fair value as the amount it would cost to pay a reinsurer to take the unpaid loss liability off your balance sheet. Two calculations are shown below for the running liability example. For these calculations, it is assumed that the immunized, risk free discount rate is 4.25% and the weighted average life of the payments is 4.5 years. The implied risk free discount factor is 0.829; discounted reserves equal $54 billion.

Fair value could be calculated using risk adjusted discount rates from Butsic [1]. Butsic shows that the appropriate risk adjusted rate, given the immunized risk free rate and the capital requirement is a liability analog to the CAPM formula.

$$r_{riskadjusted} = \left(\frac{capital}{reserves}\right)\left(ROE - r_{riskfree}\right)$$

$so,$

$$r_{riskadjusted} = \left(\frac{\$6}{\$65}\right)(20\% - 4.25\%) = 1.5\%$$

The above formula assumes a 20% pre-tax ROE target. Using the average life of about 4.5 years, the risk adjusted discount factor is 0.935 (=$1.015^{-4.5}$). Total value would be $65*0.935 = $61 billion. The implied loading for risk is $7 billion (= $61 - $54).

Alternatively, the value of the unpaid loss liabilities could be assessed using risk neutral distributions via PH-transforms as discussed by Wang [21] and Butsic [2]. The simple lognormal assumptions applied through out this paper come in handy here. A transformed lognormal entails shifting only the location parameter. If the underlying distribution has parameters μ and σ and the appropriate risk load is λ%, the risk neutral distribution will be lognormal with parameters $\mu' = \ln(1+\lambda) + \mu$, σ. A starting point for the risk load percentage could be the ratio of the risk charge, above, to required reserves shown in 3.3. Again using the industry liability data:

$$\lambda = (7.0/65) = 0.108$$

$$\mu' = \ln(1+\lambda) + \mu = \ln(1+0.108) + 4.173 = 4.275$$

$$\sigma = 0.045$$

Transforming back to dollars, the expected value of the risk neutral distribution is

$$E[X] = e^{\mu' + \frac{1}{2}\sigma^2} = \$72 \text{ billion}$$

Allowing for time value, the risk neutral value is on the order of $60 billion (= 72*.829). The results from the risk adjusted discounting methodology and the PH-transform are very similar.

One of the nice features of the PH-transform is the ability to use the risk neutral distribution to price layers of the distribution by simply taking the difference of the respective limited expected values. Perhaps this could facilitate a more active, liquid market in unpaid loss liabilities. Ironically, though, if we had a liquid market, none of the above calculations of fair value would be necessary.

Statutory risk based capital. The treatment of aggregate distributions also highlights some of the flaws with the mechanical formulas for statutory risk based capital:

1. Supporting economic capital makes sense only in the aggregate and only then when correlations have been appropriately reflected. Correlations between lines of business are imperative and cannot be ignored.

2. RBC, like economic capital, should be the difference between the aggregate value at risk (F^{-1}(1-ruin)) and the carried provision for the unpaid loss liabilities. If reserves are strengthened, required supporting capital should decrease. RBC, however, inappropriately assumes companies are currently adequately, *and only adequately,* reserved, *charging* for *any* additions to reserves. This penalizes well-reserved companies and those wishing to become better reserved.

3. Using techniques illustrated here, true economic capital requirements *can* be calculated with accepted actuarial techniques at an individual company level. Industry norms would have little use.

The above remarks refer only to RBC charges for unpaid losses. Of course other risk factors need to be integrated into a total economic capital figure. But this is only a simple extension of Wang's standard normal copula method shown above. Marginal risk distributions can be created for investment portfolios, catastrophe exposures, etc., and integrated into a total distribution for use in the calculation of required economic capital.

Best estimates. Statutory accounting requires that we establish a best estimate reserve. Further, we must establish that reserve by line of business. Further still, if management has a number of alternative estimates of unpaid loss, all equally likely, the best estimate is presumably the average of all estimates.

If regulation was principally focused on solvency, all that matters is the aggregate distribution of unpaid losses, the actual provision for unpaid losses, and the amount of capital available. In this construct, best estimates of line-by-line reserves are less relevant.

Materiality. ASOP 36 places great emphasis on materiality but does not define the term. Materiality is best defined with a distribution like that shown in Graph 2.1 (pdf) as the basis. Whether or not an issue is material depends on the answer to questions such as, "where could this issue move my estimate of unpaid losses in the *a priori* distribution of possible unpaid losses?" "Could it change the shape of the distribution?" In the end, if a company is reserved at the 51^{st} percentile (the mean of the industry aggregate distribution in our example), a material movement is one that would drop you to the N^{th} percentile in the posterior distribution.

Unfortunately, now the definition of 'material' hinges on the definition of N. Having a definition of N, however, yields an interesting implication. If a company is already just at or below the N^{th} percentile of the distribution of unpaid losses, materiality disclosures are almost a moot point. Regulatory emphasis should be placed on reserve adequacy.

Materiality is clearly a function of a company's size. A large company could conceivably have an issue looming that could move their best estimate of unpaid losses by millions of dollars, but this might be the difference of being at the 51^{st} percentile and the 50^{th}. This should not be considered material. On the other hand there are clearly companies where materiality would be measured in the thousands of dollars.

Furthermore, portfolio diversification is again key. Any actuary can think of a handful of nasty things that could cause adverse development in unpaid losses. Perhaps these nasty things could even be characterized with a mean and distribution of possible results. If we went further and thought in terms of the distributions of nasty things *and* the inter-nasty-things-correlations, we could aggregate using the same technology presented here. Would the aggregate distribution of potentially 'material' things be material? Perhaps not. Why? Because it is precisely the highly skewed, generally independent distributions that get heavily diversified away.

4. CONCLUSIONS

This paper used published and readily available data and techniques, along with a simple proposal for measuring correlations amongst unpaid losses, to produce a sample aggregate loss distribution for the U.S. commercial lines unpaid losses (1990-2000) as of year-end 2000. By the time the end of this paper mercifully came, the author had whined about a great many things. In conclusion:

1. Point estimates for unpaid loss and loss expense are insufficient. Methods to produce distributions exist and are reasonably approachable.
2. Methods exist to aggregate risk distributions, given correlations.
3. Correlations can be measured directly from the data normally employed for loss reserve analysis.
4. Aggregate distributions of unpaid losses are useful analytical tools, with implications for required economic capital, capital allocation, pricing and valuation, and various issues associated with accounting such as "best estimate," "materiality," and RBC.

ACKNOWLEDGEMENTS

The author would like to gratefully acknowledge a number people who reviewed and commented on earlier drafts of this paper. Special thanks go to Gary Venter who turned me on to the idea of using calendar year distortions to measure correlations; to Chris Gross who offered up the idea of the simplified aggregation method; Ben Zehnwirth, who, along with Chris, offered views on capital allocation; and Kyle Vrieze, who sparked the thought to include the lognormal PH-transform in the pricing example. Finally, thanks to Scott Anderson, Todd Bault, Kyle Vrieze, and Paul Wulterkens for reviewing early drafts of this work and offering their feedback. The author remains solely responsible for errors that may have persisted to the final draft, but go ahead and call one of these guys anyway if you have a complaint.

Selected Bibliography

1. Butsic, Robert P. "Determining the Proper Interest Rate for Loss Reserve Discounting: An Economic Approach," 1988 CAS Discussion Paper Program, p. 147

2. ___________. "Capital Allocation for Property-Liability Insurers: A Catastrophe Reinsurance Application," Casualty Actuarial Society Forum, Spring 1999, p. 1

3. ___________. "Report on Covariance Method for Property-Casualty Risk-Based Capital," Casualty Actuarial Society Forum, Summer 1993, p. 173

4. CAS Task Force of Fair Value Liabilities White Paper, 2000.

5. D'Arcy, Steve, et al. "Building a Public Access PC-Based DFA Model," CAS Forum, Summer 1997, Volume 2

6. Feldblum, Sholom. "NAIC Property/Casualty Insurance Company Risk-Based Capital," PCAS LXXXIII, 1996, p.297

7. _____________. "Risk Load for Insurers," PCAS LXXVII, 1990, p.160

8. Bault, Todd R. Discussion, PCAS LXXXII, 1995 p.78

9. Halliwell, Leigh. "Loss Prediction by Generalized Least Squares," PCAS LXXXIII, p. 190

10. Holmberg, Randall. "Correlation and the Measurement of the Loss Reserve Variability," Casualty Actuarial Society Forum, Spring 1994, p.247

11. Kulik, John M. "A Practical Application of Modern Portfolio Theory to Capital Allocation," Casualty Actuarial Society Forum, Spring 1999, p. 137

12. Mack, Thomas. "Distribution Free Calculation of the Standard Error of the Chain Ladder Reserve Estimates," ASTIN Bulletin, Vol. 23, No. 2, 1993, p. 213

13. Mango, Donald. "Risk Load and the Default Rate of Surplus," CAS Discussion Paper Program, May 1999, p. 175

14. Murphy, Daniel M. "Unbiased Loss Development Factors," PCAS LXXXI, 1994, p.154

15. Myers, Stewart C. and James A. Read Jr. "Line-by-line Surplus Requirements for Insurance Companies," [unpublished paper for the Automobile Insurers' Bureau of Massachusetts], April 1999

16. Nakada, Peter, et al. "P&C RAROC: A Catalyst for Improved Capital Management in the Property and Casualty Insurance Industry," The Journal of Risk Finance, Fall 1999

17. Philbrick, Stephen. "Accounting for Risk Margins," Casualty Actuarial Forum, Spring 1994, Vol. I, pp 1-87

18. Stone, James M. "A Theory of Capacity and the Insurance of Catastrophe Risks," The Journal of Risk and Insurance, June 1973, Vol. XL No. 2 (Part I) and September 1973, Vol. XL No. 3 (Part II)

19. Venter, Gary. "Testing Assumptions of Age-to-Age Factors," PCAS LXXXV, 1998, p. 807

20. Verall, Richard. "Statistical Methods for the Chain Ladder Technique," Casualty Actuarial Society Forum, Spring 1994, p.393

21. Wang, Shaun. "Implementation of Proportional Hazards Transforms in Ratemaking," PCAS LXXXV, 1998, p. 940

22. __________. "Aggregation of Correlated Risk Portfolios: Models and Algorithms," PCAS LXXXV, 1998, p. 848

23. Meyers, Glenn. Discussion, PCAS LXXXVI, 1999, p. 781

24. Zehnwirth, Ben. "Probabalistic Development Factor Models with Applications to Loss Reserve Variability, Prediction Intervals and Risk based Capital," Casualty Actuarial Society Forum, Spring 1994, p.447