The Aggregation and Correlation of Reinsurance Exposure

Glenn G. Meyers, FCAS, MAAA
Fredrick L. Klinker, FCAS, MAAA, and
David A. Lalonde, FCAS, MAAA, FCIA

The Aggregation and Correlation of Reinsurance Exposure

By

Glenn Meyers – Insurance Services Office

Fredrick Klinker – Insurance Services Office

David Lalonde – AIR Worldwide Corporation

Abstract

This paper begins with a description of how to calculate the aggregate loss distribution for a reinsurer. We include most of the standard exposures as well as property catastrophe exposure. Next we show how this aggregate loss distribution can be used to determine the needed capital, and its cost, for a reinsurer. Finally we show how to calculate the capacity charges for individual reinsurance contracts that will allow the reinsurer to recover its cost of capital. We demonstrate the use of this methodology on some illustrative reinsurance contracts. We believe this methodology can be used in practice by most reinsurers.

Authors

Glenn Meyers heads the Special Markets Products Division at Insurance Services Office. Since joining ISO in 1988 he has led the technical development of ISO's increased limits ratemaking, PSOLD – ISO's commercial property size of loss database, MILD – ISO's reinsurance exposure rating product, ISO's joint loss and ALAE distributions and ISO's Underwriting Risk Model. Glenn's academic background includes a B.S. in Mathematics and Physics from Alma College, a M.A. in Mathematics from Oakland University and a Ph.D. in Mathematics from SUNY at Albany. He is a FCAS and MAAA. He is currently serving the CAS on the Dynamic Risk Modeling Committee and the Future Education Task Force.

Fred Klinker works with Glenn in the Special Markets Products Division at Insurance Services Office. Prior to coming to ISO, he worked for seventeen years in Actuarial Research at CIGNA P&C, doing financial, mathematical, and statistical modeling. He has a BS/BA in Physics/Math from the University of Rochester and graduate work in Physics and Math from the Universities of Rochester and Pennsylvania. He is a FCAS and MAAA.

David Lalonde is Senior Vice President at AIR Worldwide Corporation responsible for the Services and Analytics group, providing Catastrophe Loss Analysis Services (CLAS™) and Risk Transfer Services (RTS™). He has been involved in the development and use of catastrophe modeling technology since 1995 and regularly assists clients in use of models for exposure management, ratemaking and underwriting, risk transfer and the securitization of insurance risk, portfolio optimization, and dynamic risk evaluation and analysis models. Mr. Lalonde is a FCAS, FCIA, MAAA. Mr. Lalonde earned his B.Math. (Honours) in Actuarial Science with Statistics from the University of Waterloo.. He is currently serving on the CAS Advisory Committee on Securitization/Risk Financing.

1. Introduction

This paper has three objectives:

1. Demonstrate a practical method to determine the distribution of a reinsurer's aggregate loss payments. This includes not only losses from the contracts it currently is reinsuring, but also contracts that have expired but still have outstanding claims. This distribution will depend on the variation of each contract's claim frequency and severity. It will also reflect dependencies among the various hazards reinsured.
2. Using the results of Objective #1, demonstrate how to determine the amount of capital needed for a reinsurance company based on its risk of loss.
3. Using the results of Objective #2 demonstrate how to determine the capacity charge for a new reinsurance contract.

We will illustrate the use of our model and methodology on a portfolio of reinsurance contracts. The parameters for the loss models were obtained from analyses by Insurance Services Office (ISO) and AIR Worldwide Corporation (AIR).

The exposures for these contracts were obtained from the annual statements for several primary insurers and from data reported to ISO. Using the descriptions of the reinsurance programs that were reported to A.M. Best Company, we modified the loss models accordingly.

We treat the time value of money by assuming a fixed risk-free interest rate. While the assets of a reinsurer are not always risk-free, a full treatment of asset risk is beyond the scope of this paper. Thus, we should expect reinsurers to have more capital than that indicated by the methodology described in this paper because they have asset risk.

We begin with a description of possible ways to model a reinsurer's distribution of underwriting losses. This description will include ways to model dependencies among the various reinsurance contracts. It will also discuss how to parameterize these models.

Next we will describe how we calculate the required capital. This description will include a short survey of the issues involved in making such a calculation. It turns out that there is no strong consensus on how to do this; but, if we are to get a final answer, we must and do pick one method.

We then move on to developing a methodology for calculating a capacity charge. As we do in our section on calculating the required capital, we will include a short survey of the issues involved in doing this. Again we note that there is no strong consensus on how to do this but, as before, we do pick one method.

While we recognize that others may differ in their methodology for solving these problems, we do feel that our methodology for calculating both the required capital and the capacity charge is reasonable. We note that the underwriting risk model that we have built to solve these problems could be used for other methodologies.

2. Models of Reinsurer Losses

This section begins with a description of the classic collective risk model, and it then enhances it with correlations or, more precisely, dependencies generated by parameter uncertainty.

Next we introduce catastrophe models, in which the dependencies are caused by geographic proximity. We describe catastrophes generated by hurricanes and earthquakes.

2.1 The Collective Risk Model

The collective risk model (CRM) describes the total reinsured loss in terms of the underlying claim count and claim severity distributions for each reinsured contract. We describe this model by the following simulation algorithm.

Simulation Algorithm #1

Step

1. For each reinsurance contract, h, with uncertain claim payments, do the following:
   - Select random claim count K_h from a distribution with mean λ_h where λ_h is the expected claim count for contract h.
   - For each h, select random claim sizes, Z_{hk}, for $k = 1,\ldots, K_h$.
2. Set $X_h = \sum_{k=1}^{K_h} Z_{hk}$ = Loss for contract h.
3. Set $X = \sum_h X_h$ = Loss for the reinsurer.

This formulation of the CRM assumes that the losses for each class are independent. We now introduce a dependency structure into the CRM with the following algorithm.

Simulation Algorithm #2

Step

1. For each reinsurance contract h, with uncertain claim payments, do the following:
   - Select a random claim count K_h from a distribution with mean λ_h where λ_h is the expected claim count for contract h.

- For each h select a random claim size, Z_{hk}, for $k = 1,\ldots,K_h$.

2. Set $X_h = \sum_{k=1}^{K_h} Z_{hk}$ = Loss for contract h.
3. Select a random β from a distribution with $E[\beta] = 1$ and $Var[\beta] = b$.
4. Set $X = \beta \cdot \sum_h X_h$ = Loss for the reinsurer.

The extra step of multiplying all the losses by a random β adds variability in a way that losses for each reinsurance contract will tend to be higher, or lower, together at the same time. This induces one kind of dependency, or correlation, among the losses of different reinsurance contracts. One can think of b as a parameter that quantifies the uncertainty in the economic environment affecting multiple lines of insurance.

Figures 1-4 provide a graphic illustration of how Simulation Algorithm #2 generates dependency and correlation. In these figures we randomly selected X_1 and X_2. Next we randomly selected β. We then plotted βX_1 against βX_2. If we do not change the distributions X_1 and X_2, a higher b will lead to a higher coefficient of correlation. But, as illustrated in Figures 3 and 4, the coefficient of correlation also depends on coefficients of variation (CV) of X_1 and X_2.

Figure 1

X_1 and X_2 are independently drawn random variables with CV=0.1.

β was drawn from a distribution with b=Var[β] = 0. Thus ρ = 0.00.

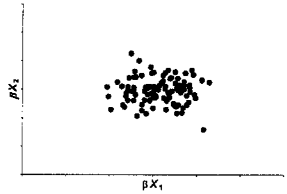

Figure 2

X_1 and X_2 are independently drawn random variables with CV=0.1

β was drawn from a distribution with b=Var[β] = 0.005. Thus ρ = 0.33.

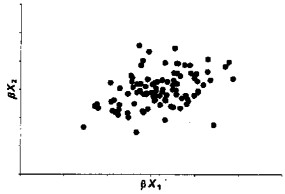

Figure 3

X_1 and X_2 are independently drawn random variables with CV=0.1

β was drawn from a distribution with b=Var[β] = 0.020. Thus ρ = 0.66.

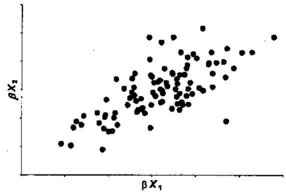

Figure 4

X_1 and X_2 are independently drawn random variables with CV=0.2

β was drawn from a distribution with b=Var[β] = 0.020. Thus ρ = 0.33.

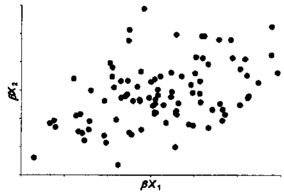

Having described one method to introduce dependencies into the collective risk model, we now apply this method to a model of the underwriting losses for a reinsurer. Here is a summary of the main features of this model.

- It is necessary to hold capital for uncertain losses in expired reinsurance contracts. Thus the model treats unpaid losses from both new and expired reinsurance contracts from prior accident years

- We use separate parameter uncertainty multipliers for both claim frequency and claim severity. For reinsurance contract h, a random claim frequency multiplier, α_h, is applied to the expected claim count parameter, λ_h. Each α_h has a mean of one and a variance of g_h. We call g_h the covariance generator for contract h.

- Each reinsurance contract is assigned to a distinct "covariance group" according to the line of business that it covers. (Granted, some reinsurance contracts cover multiple lines, but in this paper, we use a narrower definition of "contract.") Within a given covariance group, the random claim frequency multipliers, α_h, are identical within line of business, not necessarily identical to other lines of business in the same covariance group, but they increase and decrease together.

- The random claim severity multiplier, β, is applied uniformly across all contracts.

- One can informally classify the sources of risk in this model into ***process risk*** and ***parameter risk***. Process risk is the risk attributable to random claim counts and claim sizes, and parameter risk is the risk attributable to the randomness of the claim frequency multipliers and the claim severity multiplier.

- When parameter risk operates on several contracts simultaneously, we say that there is correlation generated by parameter risk.

These features are described in the following algorithm.

Simulation Algorithm #3

Step

1. Select a random β from a distribution with $E[\beta] = 1$ and $Var[\beta] = b$.
2. For each covariance group i, select random percentile p_i.
3. For each covariance group i, reinsurance contract h in the covariance group (denoted by G_i), and accident year y with uncertain claim payments, do the following:
   - Select $\alpha_{hy} = p_i^{th}$ percentile of a distribution with $E[\alpha_{hy}] = 1$ and $Var[\alpha_{hy}] = g_{hy}$
   - Select random claim count K_{hy} from a distribution with mean $\alpha_{hy} \cdot \lambda_{hy}$, where λ_{hy} is the expected claim count for reinsurance contract h and accident year y in covariance group i.
   - For each h and y, select random claim size Z_{hyk} for $k = 1, \ldots, K_{hy}$.
4. Set $X_i = \sum_{h \in G_i} \sum_{y} \sum_{k=1}^{K_{hy}} Z_{hyk} =$ Loss for covariance group i.
5. Set $X = \beta \cdot \sum_{i} X_i =$ Total loss for the reinsurer.

We now describe our parameterization of this model.

- For the non-catastrophe reinsurance contracts, we use claim severity distributions derived by ISO. We use a piecewise linear approximation to the ISO models.

- Smaller claims tend to settle quickly. In fitting the models for the distribution of future payments for expired reinsurance contracts, we removed those claims that are already settled.

- Reinsurers often write multiple contracts, covering different layers, with a single insurer. For example, one reinsurance contract will cover 50% of a lower layer, and another contract will cover 80% of a higher layer. We treat such arrangements as a single contract and adjust the claim severity distribution accordingly.

- We use the negative binomial distribution to model claim counts. The expected claim count will depend on the reinsurer's limits and exposure. A second parameter of the negative binomial distribution, called the contagion parameter must be provided. We use estimates of the contagion parameters obtained in an analysis performed by ISO. This analysis is described in the appendix.

- The same analysis in the appendix also provides estimates of the covariance generators, g_h. A noteworthy feature is that these estimates use data from several insurers. This estimation necessarily assumes that each g_h is the same for all insurers writing that particular line of insurance. While we agree in principle that each g_h could differ by insurer, it is unlikely that any single insurer will have enough observations to get reliable estimates of the g_h's.

- The main idea behind the estimation of the parameters, described in the appendix, is that expected values of various statistics that we can calculate from the data are functions of the negative binomial parameters and the covariance generators. We calculated these statistics for a large number of insurance companies and we found parameter values that best fit the statistics we calculated. As we show in the appendix, reliable estimates of these parameters cannot be obtained with data from a single insurer. It is only by combining the data of several insurers that we can obtain reliable estimates of these parameters.

Finally, we describe how we calculate a reinsurer's distribution of underwriting losses. Since we describe the loss model in terms of a computer simulation, one could actually do the simulations. In practice, many do. We calculate the distribution of underwriting losses with Fourier transforms using the method described by Heckman and Meyers [1983]. The extension of this method to address dependencies is described by Meyers [1999a and 1999b].

Both simulation and Fourier transforms are valid ways to calculate the distribution of underwriting losses. The advantage of Fourier transforms is that one can calculate the distribution of underwriting losses in seconds, where a simulation could take hours to do the same task. A disadvantage of Fourier transforms is that it can take a long time to do the initial set-up whereas the set-up time for a simulation is relatively short.

2.2 Catastrophic Perils

Natural catastrophes such as earthquakes, hurricanes, tornadoes, and floods have an impact on many insureds; and the accumulation of losses to an insurer or reinsurer can jeopardize the financial well-being of an otherwise stable, profitable company. Hurricane Andrew, in addition to causing more than $16 billion in insured damage, left at least 11 companies insolvent in 1992. The 1994 Northridge earthquake caused more than $12 billion in insured damage in less than 60 seconds.

Fortunately, such events are infrequent. But it is exactly their infrequency that makes the estimation of losses from future catastrophes so difficult. The scarcity of historical loss data makes standard actuarial techniques of loss estimation inappropriate for quantifying catastrophe losses. Furthermore, the usefulness of the loss data that does exist is limited because of the constantly changing landscape of insured properties. Property values change, building codes are change over time, along with the costs of repair and replacement. Building materials and designs change, and new structures may be more or less vulnerable to catastrophic events than were the old ones. New properties continue to be built in areas of high hazard. Therefore, the limited loss information that is available is not sufficient for directly estimating future losses.

The modeling of catastrophes is based on sophisticated stochastic simulation procedures and powerful computer models of how natural catastrophes behave and act upon the man-made environment. The modeling is broken into four components. The first two components, event generation and local intensity calculation, define the hazard. The interaction of the local intensity of an event with specific exposures is developed through engineering based vulnerability functions in the damage estimation component. In the final component, insured loss calculation, policy conditions are applied to generate the insured loss.

Figure 5 below illustrates the component parts of the AIR state-of-the-art catastrophe models. It is important to recognize that each component, or module, represents both the analytical work of the research scientists and engineers who are responsible for its design and the complex computer programs that run the simulations.

Figure 5: Catastrophe Model Components (in gray)

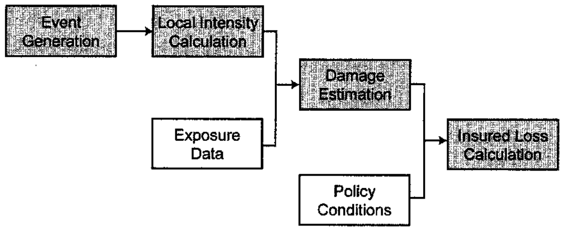

2.2a Event Generation Module

The event generation module determines the frequency, magnitude, and other characteristics of potential catastrophe events by geographic location. This requires, among other things, a thorough analysis of the characteristics of historical events.

After rigorous data analysis, researchers develop probability distributions for each of the variables, testing them for goodness-of-fit and robustness. The selection and subsequent refinement of these distributions are based not only on the expert application of statistical techniques, but also on well-established scientific principles and an understanding of how catastrophic events behave.

These probability distributions are then used to produce a large catalog of simulated events. By sampling from these distributions, the model generates simulated "years" of event activity. Many thousands of these scenario years are generated to produce the complete and stable range of potential annual experience of catastrophe event activity and to ensure full coverage of extreme (or "tail") events, as well as full spatial coverage.

2.2.b Local Intensity Module

Once the model probabilistically generates the characteristics of a simulated event, it propagates the event across the affected area. For each location within the affected area, local intensity is estimated. This requires, among other things, a thorough knowledge of the geological and/or topographical features of a region and an understanding of how these features are likely to influence the behavior of a catastrophic event. The intensity experienced at each site is a function of the magnitude of the event, distance from the source of the event, and a variety of local conditions. Researchers base their calculations of local intensity on empirical observation as well as on theoretical relationships between the variables.

2.2.c Damage Module

Scientists and engineers have developed mathematical functions called damageability relationships, which describe the interaction between buildings (both their structural and nonstructural components as well as their contents) and the local intensity to which they are exposed. Damageability functions have also been developed for estimating time element losses. These functions relate the mean damage level as well as the variability of damage to the measure of intensity at each location. Because different structural types will experience different degrees of damage, the damageability relationships vary according to construction materials and occupancy. The model estimates a complete distribution around the mean level of damage for each local intensity and each structural type and, from there, constructs an entire family of probability distributions. Losses are calculated by applying the appropriate damage function to the replacement value of the insured property.

The AIR damageability relationships incorporate the results of well-documented engineering studies, tests, and structural calculations. They also reflect the relative effectiveness and enforcement of local building codes. Engineers refine and validate these functions through the use of post-disaster field survey data and through an exhaustive analysis of detailed loss data from actual events.

2.2.d Insured Loss Module

In this last component of the catastrophe model, insured losses are calculated by applying the policy conditions to the total damage estimates. Policy conditions may include deductibles by coverage, site-specific or blanket deductibles, coverage limits and sublimits, loss triggers, coinsurance, attachment points and limits for single or multiple location policies, and risk-specific reinsurance terms.

2.2.e Model Output

After all of the insured loss estimations have been completed, they can be analyzed in ways of interest to risk management professionals. For example, the model produces complete probability distributions of losses, also known as exceedance probability curves (see Figure 6). Output includes probability distributions of gross and net losses for both annual aggregate and annual occurrence losses. The probabilities can also be expressed as return periods. That is, the loss associated with a return period of 10 years is likely to be exceeded only 10 percent of the time or, on average, in one year out of ten. For example, the model may indicate that, for a given regional book of business, $70 million or more in insured losses would be expected to result once in 50 years, on average, in a defined geographical area, and that losses of $175 million or more would be expected, on average, once every 250 years.

Output may be customized to any desired degree of geographical resolution down to location level, as well as by line of business and, within line of business, by construction class, coverage, etc. The model also provides summary reports of exposures, comparisons of exposures and losses by geographical area, and detailed information on potential large losses caused by extreme "tail" events.

Figure 6: Exceedance Probability Curve (Occurrence)

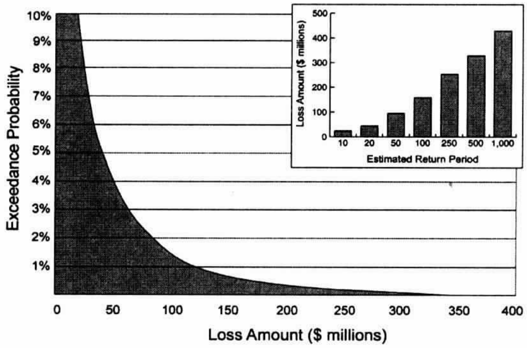

2.2.f Correlation

An advantage of this modeling approach is the generation of a stochastic event set that can be used to analyze multiple exposure sets. In this study, individual companies' exposures were analyzed using a common catalog of events. As mentioned earlier, details of reinsurance programs were also applied, resulting in both net and gross distributions of potential catastrophe losses. By analyzing various sets of exposure against the same set of events we are able to ascertain correlation amongst the exposure sets.

3. Calculating the Required Capital

This paper is focused on the underwriting risk generated by uncertain loss payments. We assume that all assets are invested at a risk-free rate of return and thus make the simplifying assumption that the capital required by a reinsurer depends solely on its aggregate loss distribution.

A reinsurer is exposed to underwriting risk not only from future claims on new business, but also from unsettled claims on past business. One must consider the underwriting risk from both sources when calculating the required capital. Larger claims tend to take longer to settle, and the underwriting loss model should reflect this.

Let X be the random variable for the reinsurer's total loss. Denote by $\rho(X)$ the total assets that the reinsurer needs to support its business[1]. Now some of the reinsurer's assets come from the premium it charges for its business. At a minimum, this amount should equal the expected value of X, $E[X]$. The remaining assets, which we call (economic) capital, must come from investors. We define the capital needed by the reinsurer by the equation:

$$\text{Capital} = \rho(X) - E[X] \qquad (1)$$

Let α be a selected percentile of X. The tail value-at-risk for X, $TVAR_\alpha(X)$, is defined to be the average of all losses greater than or equal to the α^{th} percentile of X. In this paper we use $\rho(X) = TVAR_{99\%}(X)$.

[1] If we were to allow assets, denoted by A, to be random, we would require A to satisfy $\rho(X–A) = 0$. With translation invariance, this says that $\rho(X) = A$ when A is fixed.

The tail value-at-risk is a member of an important class of risk measures, called coherent measures of risk. These measures are defined by the following set of axioms.

1. Subadditivity — For all random losses X and Y,

$$\rho\,(X+Y) \leq \rho\,(X) + \rho\,(Y).$$

2. Monotonicity — For all random losses X and Y, if $X \leq Y$ for all scenarios, then

$$\rho\,(X) \leq \rho\,(Y).$$

3. Positive Homogeneity — For all $\lambda \geq 0$ and random losses X,

$$\rho\,(\lambda X) = \lambda\rho\,(X).$$

4. Translation Invariance — For all random losses X and constant loss amounts α,

$$\rho\,(X+\alpha) = \rho\,(X) + \alpha.$$

These measures were originated by Artzner, *et al.* [1999]. See Meyers [2002] for an elementary description of these measures as well as for other coherent measures of risk.

4. Calculating the Capacity Charge

As noted in the last section, a reinsurer needs to get capital from investors in order to attract business. The investors expect to be compensated in return for providing this capital at an expected rate of return that is somewhat higher than they would obtain for not exposing their capital to reinsurance risk. This additional return must come from the sum of the premiums charged to each individual reinsurance contract. The portion of this additional return for an individual reinsurance contract is called the capacity charge. In this section we give our formula for calculating the capacity charge.

Our formula is based on the underwriting strategy of establishing a target return on the additional capital needed to write this contract. We view the capacity charge as input into the underwriting decision. If the market will not allow the reinsurer to obtain this target return, the reinsurer should consider not writing the proposed contract.

We divide this section into two subsections. The first subsection gives our rationale for using this formula in terms of our chosen underwriting strategy. The second subsection gives our capacity charge formula.

4.1 A General Discussion of Capacity Charge Formulas

We take it as a given that a sound method of calculating capacity charges should lead to decisions that benefit the entire operation of a reinsurer.

This discussion will be somewhat informal. A more rigorous treatment of this subject is provided by Meyers [2003]. We shall quote a number of results that are proved in that paper.

Proposition 1

Adding a reinsurance contract to a reinsurer's portfolio will increase the reinsurer's expected return on capital if and only if the contract's expected return on marginal capital (i.e., the contract's capacity charge divided by the additional capital needed to write the contract) is greater than the reinsurer's current expected return on capital.

This proposition provides a minimum standard on the capacity charge for *new* contracts. It says nothing about the capacity charge on existing contracts.

Proposition 2

Let the reinsurer's capital be determined by Equation (1), with $\rho(X)$ being a subadditive measure of risk. Then the sum of the marginal capitals for each reinsurance contract is less than or equal to the reinsurer's total capital.

As we shall see in the examples below, we expect strict inequality to be common. When this is the case, at least some of the contracts will have an expected return on marginal capital that is greater than the reinsurer's overall return on capital. However there are conditions when we can prove that the sum of the marginal capitals will be equal to the total capital.

Definition 1

Suppose for a reinsurance contract i, the random losses, X_i, for the contract are equal to a random number, U_i, times the exposure measure, e_i, for all possible values of e_i. Then, following Mildenhall [2002], the distribution of X_i is said to be ***homogeneous*** with respect to the exposure measure, e_i.

Proposition 3

Assume that the needed capital is a smooth (differentiable) function of the exposure.

Let the random loss, X_i, for the i^{th} reinsurance contract be a homogeneous random variable for each contract with respect to some exposure measure, e_i.

Let $X = \sum_i X_i$. Let the reinsurer's capital be determined by Equation (1), with $\rho(X)$ being a measure of risk satisfying the positive homogeneity axiom. Then the sum of the marginal capitals for each reinsurance contract is equal to the reinsurer's total capital.

An early version of Proposition 3, assuming each X_i has a lognormal distribution and using a different formula for calculating the needed capital, was proved by Myers and Read [2001]. Mildenhall [2002] proved that the homogeneity assumption was both necessary and sufficient for the Myers-Read result. The proof of Proposition 3 above is a direct consequence of Lemma 2 in Mildenhall's paper.

Note that the definition of homogeneity bears a strong resemblance to the way we introduce parameter risk in Section 2 above. As the exposure (in Section 2, quantified by the expected claim counts λ_{hy}) increases, the parameter risk becomes an increasingly large part of the total risk. But in the parameterization of our model, the parameter risk is rarely dominant enough to assume homogeneity.

Proposition 4

Assume that the needed capital is a smooth (differentiable) function of the exposure. If we can continuously adjust the exposures while holding the needed capital constant, the maximum expected return on capital occurs when the expected return on marginal capital is the same for all contracts.

Note that Proposition 4 does not require homogeneity with respect to some measure of exposure. If the loss random variables are not homogeneous, the equal expected returns on marginal capital under the optimality conditions of the proposition will be higher than the reinsurer's overall return on capital.

Definition 2

The ***heterogeneity multiplier, HM,*** for a reinsurer is its needed capital divided by the sum of the marginal capitals for each contract in its reinsurance portfolio.

The motivation for this definition arises from the fact that most reinsurers will have a total capital that is higher than the sum of the marginal capitals for each reinsurance contract. In theory, a market could evolve with bigger contracts where parameter risk dominates the process risk, and the homogeneity conditions required by Proposition 3 would be reasonable. In practice, the distribution of losses of reinsurance contracts are far from homogeneous, and the heterogeneity multiplier for a given reinsurer will be noticeably higher than the theoretical minimum of 1.

Our target capacity charge will be determined by a target return on marginal capital times the reinsurer's heterogeneity multiplier. To summarize, the rationale for this is based on:

1. Proposition 4 – The expected return on marginal capital should be equal for all contracts to if the reinsurer is to make the most efficient use of its capital.
2. Propositions 2 and 3 – The sum of the marginal capitals over all reinsurance contracts is less than or equal to the total capital. The conditions that will force equality are not satisfied.

Note that the rationale underlying this charge depends on the individual reinsurance contracts being a small part of a reinsurer's portfolio, so that the smoothness criterion of Proposition 3 and 4 is a reasonable assumption.

4.2 The Capacity Charge Formula

If the underwriting result of all reinsurance contracts could be known within a year, we expect to release the capital at the end of the year, earning investment income on the capital. We would calculate the capacity charge as follows.

1. Establish a reference portfolio, Π, of existing contracts. Ideally this portfolio is updated as each new contract is accepted. But the year-end rush to book January 1 renewals makes this difficult to do in practice, so the portfolio will be set up according to a business plan.

2. Calculate the marginal capital for each contract in the reference portfolio. This is done by first calculating the capital needed for the reference portfolio according to Equation 1. Next we calculate the capital needed when a given contract is removed from the portfolio. The marginal capital for the given contract is the difference between the two capital calculations.

3. Calculate the heterogeneity multiplier, *HM*, by the formula:

$$HM = \frac{\text{Total capital for } \Pi}{\text{Sum of the marginal capitals of each contract in } \Pi}. \tag{2}$$

4. For a prospective reinsurance contract, calculate the marginal capital, ΔC, needed when the contract is added to the reference portfolio.

5. Let r be the rate of return needed to attract the needed capital. Let i be the rate of return on invested assets. We expect $r > i$. The capacity charge, ΔP, for the prospective reinsurance contract is given by[2]:

$$\Delta P = \frac{(r-i)\cdot HM \cdot \Delta C}{(1+r)}. \tag{3}$$

Because of the way we defined the *HM*, the sum of the capacity charges will yield the reinsurer's desired return on its capital.

[2] This formula is special case of Equation 5, which is derived below.

As discussed above, the underwriting result of some reinsurance contracts can be uncertain for a period of several years. In this case the reinsurer must hold capital over this period to support these potential liabilities. This affects the calculation of the capacity charge in the following ways.

- The reference portfolio must contain the contracts that have expired but still have uncertain losses. The required capital for the reference portfolio must reflect the uncertainty in the ultimate losses from these unexpired contracts.
- When calculating the capacity charge, the reinsurer has to consider the fact that it has to hold additional capital in future years to support the contracts it is writing now. The cost of holding this capital over this extended period of time must be included in the capacity charge.

With these considerations in mind, we calculate the capacity charge as follows.

1. Establish a reference portfolio of existing contracts for this year, and for as many years in the future that the reinsurer expects uncertainty in its ultimate losses for contract written in this and prior years. Denote the portfolio for the current year by Π_0, the portfolio for next year by Π_1, and so on. These reference portfolios will contain current and expired and planned future reinsurance contracts.
2. Calculate the marginal capital for each reinsurance contract (current and expired) in each of portfolios, Π_0, Π_1....
3. For each portfolio, Π_n, calculate the heterogeneity multiplier, HM_n, by the formula.

$$HM_n = \frac{\text{Total capital for } \Pi_n}{\text{Sum of the marginal capitals of each contract in } \Pi_n}. \quad (4)$$

4. For a prospective reinsurance contract, calculate the marginal capital, ΔC_n, needed when the contract is added to the n^{th} reference portfolio. Note that the contract is considered to be new in Π_0, but expired in Π_n, for $n > 0$.

5. Let r be the rate of return needed to attract the needed capital. Let i be the rate of return on invested assets. The capacity charge, ΔP, for the prospective reinsurance contract is given by:

$$\Delta P = \sum_{n=0}^{\infty} \frac{(r-i) \cdot HM_n \cdot \Delta C_n}{(1+r)^{n+1}}. \quad (5)$$

Note that the capacity charge is applied to the new contracts only. The time to collect the capacity charge on the expired contracts was when the contract was written. In defining the capacity charge in this way, the reinsurer is making its desired rate of return on its allocated cost of capital[3].

We finish this section with the derivation of Equation 5.

- The reinsurer puts up an initial investment of $HM_0 \cdot \Delta C_0$ at time $t = 0$. It commits to holding $PRM_1 \cdot \Delta C_1$ at time $t = 1$, $HM_2 \cdot \Delta C_2$ at time $t = 2$, and so on.
- While the reinsurer is holding the capital, it is earning interest at rate i. At time $t = 1$, it expects to receive $HM_0 \cdot \Delta C_0 \cdot (1 + i) - HM_1 \cdot \Delta C_1$. At time $t = 2$, it expects to receive $HM_1 \cdot \Delta C_1 \cdot (1 + i) - HM_2 \cdot \Delta C_2$, and so on.
- The capacity charge is equal to the initial investment less the present value (at interest rate r) of the expected amount received. That is:

[3] One possible enhancement of Equation 5 would be to vary the rate of return by the length of time it is invested.

$$\Delta P = HM_0 \cdot \Delta C_0 - \sum_{n=0}^{\infty} \frac{HM_n \cdot \Delta C_n \cdot (1+i) - HM_{n+1} \cdot \Delta C_{n+1}}{(1+r)^{n+1}}. \qquad (6)$$

- Equation 5 is derived from Equation 6 by rearranging and grouping the terms in increasing order of n.

5. Examples

We now illustrate the use of our model and methodology on a number of sample reinsurance contracts. We constructed a reference portfolio of reinsurance contracts from real insurance companies, based on publicly available data. The lines of business in the reference portfolio included general liability, commercial auto, workers' compensation, professional liability, commercial multi-peril, fire, allied lines, earthquake and some personal lines. We treated the hurricane and earthquake exposures as separate contracts that took a 25% share of the underlying catastrophe reinsurance contract.

What follows is a description of the steps we took to construct this reference portfolio.

1. We first estimated the expected direct losses, by annual statement line of business, for the insurers included in the reference portfolio. For the most recent accident year, we estimated the expected losses by multiplying the reported premiums by our estimated loss ratio for the industry. For prior accident years, we used the insurers' reported loss reserves.

2. Using details of each insurer's reinsurance program reported to the A.M. Best Company, and the loss distributions underlying the ISO Underwriting Risk Model and the AIR catastrophe model for hurricanes and earthquakes, we partitioned the insurer's expected direct losses into two segments by

annual statement line of business – the expected net losses and the expected reinsured losses.

3. In order to project the expected release of the marginal capital over the next several years using Equation 5, we need to know the marginal capital attributed to current contracts in future accident years. This will depend on the reinsurer's business plan. We assumed that the reinsurer would continue its current business plan but, going forward, we estimated the expected unpaid losses using an ISO industry loss reserve study.

For the reference portfolio, the total expected loss for all lines in the current year is $739,998,127. The expected payout for losses from prior accident years is $1,813,101,644. In constructing this reference portfolio, we did not have all the detailed contract level information that is potentially available to a reinsurer. The reference portfolio had a few hundred contracts covering a variety of limits. The insurers in the reference portfolio tended to be larger than average and thus we expect the size and the retentions of the contracts to be a bit higher than normal.

We now describe how we calculated the necessary capital for the reference portfolio, with and without the proposed contracts.

1. In evaluating the non-catastrophe exposure we used the expected loss estimates and the limits for each contract. Using claim severity distributions in the ISO Underwriting Risk Model, we obtained the expected claim count by dividing the expected loss by the expected claim severity. The claim count distribution requires a second parameter that ISO obtained from analyses similar to that described in the appendix.

2. Using exposures that primary insurers reported to ISO, we ran the AIR catastrophe model to produce 10,000 simulated years of hurricane and

earthquake losses for each primary insurer in the reference portfolio. The catastrophe losses were adjusted to reflect the reinsurance provisions and the 25% share of the catastrophe contracts taken by the reference portfolio. The losses for all the catastrophe contracts in the reference portfolio were summed by year to produce a combined catastrophe size of loss distribution.

3. The distributional information above was used to derive the reference and marginal aggregate loss distributions by a procedure mathematically equivalent to Simulation Algorithm 3 above. Table 1 describes the aggregate loss distribution for the reference portfolio.

4. Following Equation 1, we set the needed capital for the reference portfolio equal to $TVaR_{99\%}(X) - E[X]$ = \$670,997,012.

The next step was to calculate the heterogeneity multiplier, HM_n, for each year. This is done by finding the marginal capital for each reinsurance contract in the reference portfolio and applying Equation 2. While the heterogeneity multipliers varied slightly by year, they were all close to 1.64, well above the theoretical minimum of 1.00. Since we were assuming a stable business plan, we selected $HM_n = 1.64$ for all n.

Now we are ready to calculate the capacity charges for prospective reinsurance contracts using Equation 5.

The first set of examples consists of some standard property and casualty reinsurance contracts. We first calculate the marginal capital for the prospective contract for the current year and up to the following six years, which we assumed will have uncertainty in the ultimate paid losses. In this example, we are ignoring all uncertainty in ultimate losses after two years for Fire, after five years for

Commercial Auto, and after seven years for General Liability. In Table 2 we provide an illustrative aggregate loss distribution when a General Liability reinsurance contract are added to the reference portfolio. Table 3 gives the result of marginal capital calculations for the remaining contracts in this set of examples.

We used $HM_n = 1.64$, $r = 18\%$ and $i = 6\%$. The capacity charges calculated using Equation 5 for this set of examples are in Table 4.

Table 1

Aggregate Loss Distribution
Produced by
ISO Underwriting Risk Model
Reference Portfolio – Year 1

Aggregate Mean	2,553,099,771
Aggregate Std. Dev	226,983,918

Aggregate Loss	Cumulative Probability	Tail Value at Risk	Implied Capital
2,544,328,941	0.50000	2,733,373,669	180,273,898
2,572,708,866	0.55000	2,752,806,019	199,706,248
2,601,822,082	0.60000	2,773,508,440	220,408,669
2,632,192,534	0.65000	2,795,884,342	242,784,571
2,664,579,556	0.70000	2,820,498,110	267,398,339
2,699,943,781	0.75000	2,848,200,965	295,101,194
2,739,710,696	0.80000	2,880,411,932	327,312,161
2,787,036,572	0.85000	2,919,645,004	366,545,233
2,847,436,074	0.90000	2,971,590,416	418,490,645
2,887,426,613	0.92500	3,006,495,925	453,396,154
2,940,100,948	0.95000	3,053,590,102	500,490,331
2,953,219,034	0.95500	3,065,480,084	512,380,313
2,967,653,154	0.96000	3,078,622,916	525,523,146
2,983,799,371	0.96500	3,093,339,811	540,240,040
3,002,061,116	0.97000	3,110,102,957	557,003,187
3,023,031,771	0.97500	3,129,649,532	576,549,761
3,048,079,271	0.98000	3,153,220,036	600,120,266
3,080,209,479	0.98500	3,183,009,313	629,909,542
3,123,377,033	0.99000	3,224,096,783	670,997,012
3,195,198,671	0.99500	3,292,456,190	739,356,419
3,350,378,069	0.99900	3,446,040,482	892,940,711
3,416,123,232	0.99950	3,512,255,729	959,155,958
3,567,277,277	0.99990	3,670,239,190	1,117,139,419

Table 2

Aggregate Loss Distribution
Produced by
ISO Underwriting Risk Model
Reference Portfolio + General Liability Treaty B – Year 1

Aggregate Mean	2,554,099,777
Aggregate Std. Dev	227,010,259

Aggregate Loss	Cumulative Probability	Tail Value at Risk	Implied Capital
2,545,328,752	0.50000	2,734,394,860	180,295,083
2,573,712,090	0.55000	2,753,829,391	199,729,614
2,602,827,651	0.60000	2,774,534,227	220,434,450
2,633,202,873	0.65000	2,796,912,584	242,812,807
2,665,594,329	0.70000	2,821,528,848	267,429,071
2,700,955,923	0.75000	2,849,235,353	295,135,576
2,740,735,861	0.80000	2,881,449,351	327,349,574
2,788,066,006	0.85000	2,920,686,432	366,586,655
2,848,476,835	0.90000	2,972,637,184	418,537,407
2,888,468,111	0.92500	3,007,545,936	453,446,159
2,941,144,453	0.95000	3,054,644,769	500,544,992
2,954,263,865	0.95500	3,066,535,955	512,436,178
2,968,702,964	0.96000	3,079,679,982	525,580,205
2,984,851,364	0.96500	3,094,397,423	540,297,646
3,003,106,105	0.97000	3,111,162,780	557,063,003
3,024,076,687	0.97500	3,130,710,903	576,611,126
3,049,146,274	0.98000	3,154,284,003	600,184,226
3,081,283,973	0.98500	3,184,073,593	629,973,816
3,124,409,548	0.99000	3,225,162,418	671,062,641
3,196,271,436	0.99500	3,293,544,811	739,445,034
3,351,471,608	0.99900	3,447,105,396	893,005,619
3,417,154,542	0.99950	3,513,336,592	959,236,815
3,568,402,012	0.99990	3,671,307,958	1,117,208,181

Table 3

Reinsurance Contract	Marginal Capital Needed at Beginning of Year Year 0	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6
Fire A	52,488	11,869					
Fire B	54,694	12,428					
Fire C	66,358	15,383					
Comm Auto Liab A	34,962	28,845	18,913	16,435	7,944		
Comm Auto Liab B	37,350	30,810	18,045	15,255	7,533		
Comm Auto Liab C	52,799	44,260	27,308	19,896	9,560		
Comm Auto Liab D	40,810	33,850	21,319	16,976	8,064		
General Liability A	63,628	53,837	44,341	38,707	22,441	15,493	12,034
General Liability B	65,629	55,336	45,939	39,968	24,076	17,144	13,550
General Liability C	77,826	65,733	55,518	49,518	33,768	25,682	20,273
General Liability D	67,488	56,882	47,205	41,727	25,945	18,759	14,742

Table 4

Contract	First Contract Retention	First Contract Limit	Second Contract Retention	Second Contract Limit	Expected Loss	Capacity Charge	Cap Chg as % Exp Loss
Fire A	500,000	500,000	–	–	1,000,000	10,432	1.04%
Fire B	1,000,000	1,000,000	–	–	1,000,000	10,878	1.09%
Fire C	1,000,000	5,000,000	–	–	1,000,000	13,241	1.32%
Comm Auto Liab A	500,000	500,000	–	–	1,000,000	14,525	1.45%
Comm Auto Liab B	1,000,000	1,000,000	–	–	1,000,000	14,942	1.49%
Comm Auto Liab C	1,000,000	5,000,000	–	–	1,000,000	21,174	2.12%
Comm Auto Liab D	500,000	500,000	2,000,000	2,000,000	1,000,000	16,561	1.66%
General Liability A	500,000	500,000	–	–	1,000,000	31,265	3.13%
General Liability B	1,000,000	1,000,000	–	–	1,000,000	32,484	3.25%
General Liability C	1,000,000	5,000,000	–	–	1,000,000	39,976	4.00%
General Liability D	500,000	500,000	2,000,000	2,000,000	1,000,000	33,695	3.37%

When you examine Tables 3 and 4, we hope you would agree that the capacity charges follow a logical progression in terms of relative risk and the length of time that capital must be held to support that risk.

Next we consider a set of examples consisting of catastrophe treaties. As we did in constructing the reference portfolio, using exposures that primary insurers reported to ISO, we ran the AIR catastrophe model to produce 10,000 simulated years of hurricane and earthquake losses for a number of primary insurers. The catastrophe losses were adjusted to reflect the reinsurance provisions, and we continued with the 25% quota share provision that was taken by the reference portfolio. For each contract, the losses were added to the losses of the catastrophe contracts in the reference portfolio by year to produce a combined catastrophe size of loss distribution. We assumed that there was no uncertainty in the catastrophe losses after one year.

Table 5 gives an illustrative aggregate loss distribution when a catastrophe contract is added to the reference portfolio. Table 6 gives the results of the marginal capital calculations and resulting capacity charges for each of the catastrophe contracts.

Table 5

Aggregate Loss Distribution
Produced by
ISO Underwriting Risk Model
Reference Portfolio + Earthquake C

Aggregate Mean	2,559,254,438
Aggregate Std. Dev	230,864,879

Aggregate Loss	Cumulative Probability	Tail Value at Risk	Implied Capital
2,549,206,753	0.50000	2,742,159,379	182,904,941
2,577,918,536	0.55000	2,762,007,818	202,753,380
2,607,399,694	0.60000	2,783,186,946	223,932,508
2,638,206,921	0.65000	2,806,117,885	246,863,447
2,671,104,813	0.70000	2,831,392,904	272,138,466
2,707,046,998	0.75000	2,859,915,077	300,660,639
2,747,671,983	0.80000	2,893,170,794	333,916,356
2,796,075,394	0.85000	2,933,840,121	374,585,683
2,858,223,270	0.90000	2,987,982,052	428,727,614
2,899,467,753	0.92500	3,024,575,910	465,321,472
2,954,132,749	0.95000	3,074,238,226	514,983,788
2,967,817,054	0.95500	3,086,832,661	527,578,223
2,982,939,047	0.96000	3,100,778,128	541,523,689
2,999,816,143	0.96500	3,116,425,597	557,171,159
3,018,878,770	0.97000	3,134,301,652	575,047,214
3,040,922,350	0.97500	3,155,211,119	595,956,681
3,067,778,750	0.98000	3,180,494,563	621,240,125
3,101,546,799	0.98500	3,212,599,452	653,345,014
3,146,195,429	0.99000	3,257,314,564	698,060,126
3,224,424,972	0.99500	3,331,993,080	772,738,642
3,396,310,134	0.99900	3,503,462,403	944,207,965
3,468,506,073	0.99950	3,578,282,216	1,019,027,778
3,642,049,455	0.99990	3,762,498,626	1,203,244,188

Table 6

Reinsurance Contract	Marginal Capital	Expected Loss	Capacity Charge	Cap Chg as % Exp Loss
Earthquake A	14,736	5,287	2,458	46.48%
Earthquake B	7,538,096	4,939,820	1,257,201	25.45%
Earthquake C	27,063,114	6,154,667	4,513,577	73.34%
Earthquake D	1,483,536	1,273,219	247,424	19.43%
Earthquake E	1,862,063	303,947	310,554	102.17%
Earthquake F	3,174,465	593,735	529,436	89.17%
Earthquake G	5,513,907	2,760,151	919,608	33.32%
Earthquake H	2,102,509	371,200	350,656	94.47%
Hurricane A	2,092,047	123,008	348,911	283.65%
Hurricane B	95,297	75,723	15,894	20.99%
Hurricane C	3,532,354	640,824	589,125	91.93%
Hurricane D	1,838,135	462,064	306,564	66.35%
Hurricane E	1,716,522	266,411	286,281	107.46%
Hurricane F	4,063,674	226,776	677,738	298.86%
Hurricane G	2,871,555	577,426	478,917	82.94%
Hurricane H	33,428,704	7,840,572	5,575,228	71.11%
Hurricane I	22,259,834	2,197,780	3,712,488	168.92%
Hurricane J	8,167,187	2,695,188	1,362,121	50.54%

A noteworthy feature of this last set of examples is the wide range of capacity charges. Earthquake Contracts E and G provide one of the nicer illustrations of what drives these differences. Table 7 gives some key statistics.

Table 7

Earthquake Contract	**Coefficient of Variation**	**Correlation with Reference Portfolio**
E	36.3	0.76
G	27.5	0.16

Earthquake Contract E is more volatile than Earthquake G, but the main difference is the correlation between the contracts and the reference portfolio. Figure 7 provides scatter plots of the contracts and the reference portfolio.

Figure 7

Note: The vertical and horizontal dotted lines represent the respective 99th percentile of the reference portfolio and the indicated reinsurance contract for earthquake reinsurance.

Contract E

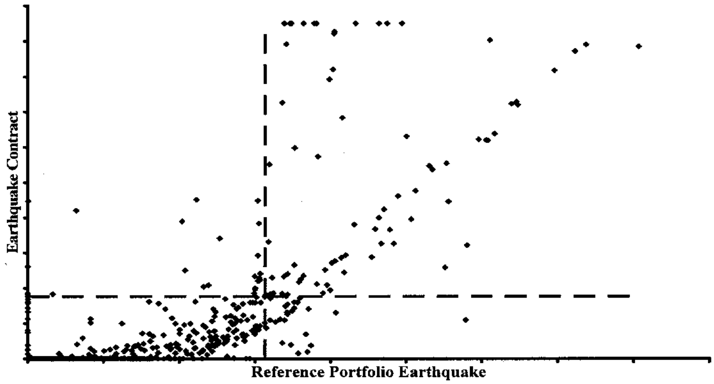

Contract G

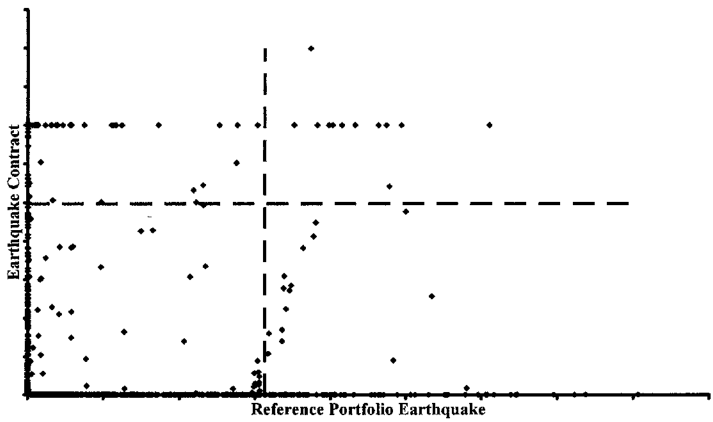

7. Summary and Conclusions

This paper started with three objectives:

1. Demonstrate a practical method to determine the distribution of a reinsurer's aggregate loss payments.
2. Using the results of Objective #1, demonstrate how to determine the amount of capital needed for a reinsurance company based on its risk of loss.
3. Using the results of Objective #2 demonstrate how to determine the capacity charge for a new reinsurance contract.

We demonstrated our methodology for accomplishing these objectives on an illustrative reinsurer with hundreds of reinsurance contracts.

We used the ISO Underwriting Risk Model to determine the aggregate loss distribution. As input, the model took the limits and quota share percentages for each reinsurance contract for the "standard lines" of insurance. We used the claim count and claim severity distributions provided by the model. For hurricane and earthquake losses, we used the AIR catastrophe model with exposures reported to ISO as input.

Dependencies among the various lines of insurance were reflected in the model by quantifications of parameter uncertainty in the standard lines of insurance and by geographic proximity for the catastrophe exposure.

Next we determined the capital needed for the reinsurer by calculating the Tail Value-at-Risk from the aggregate loss distribution.

Finally we calculated capacity charges for a variety of reinsurance contracts. The rationale underlying these calculations was that the total capacity charge over all reinsurance contracts should provide the reinsurer with a competitive expected return on capital.

The underwriting strategy used to get this expected return assumed that the reinsurer will write those contracts that provide the greatest return on marginal capital. Now it can take several years for some reinsurance contracts to be settled. The reinsurer must hold capital as long as there is uncertainty in the final settlement of its claims, and the capacity charge reflects how long capital must be held because it reinsures a given contract.

We believe we have demonstrated that this methodology can be implemented for most reinsurers.

8. Additional Comment

There is recent actuarial literature on "correlation in the tails" such as that of Venter [2002]. The analysis documented in the appendix of this paper estimates an overall level of correlation not attributed to particular region of the loss distribution. We doubt that we have sufficient data to make such an attribution.

Furthermore, to the extent that correlation in the tails is driven by large natural catastrophes, we argue that, when we couple a collective risk model parameterized by the parameters estimated in the appendix with simulation runs from a catastrophe model, as documented above, we do indeed capture at least some "correlation in the tails."

Should a reinsurer want to use a copula, or some other dependency model, our methodology for determining the needed capital and capacity charges can accommodate it. At the very least, one can generate a large number of stochastic scenarios and incorporate that into the collective risk model in exactly the same way that we did for the catastrophe model.

References

1. Philippe Artzner, Freddy Delbaen, Jean-Marc Eber and David Heath, "Coherent Measures of Risk," Math. Finance 9 (1999), no. 3, 203-228 http://www.math.ethz.ch/~delbaen/ftp/preprints/CoherentMF.pdf

2. Philip E. Heckman and Glenn G. Meyers, "The Calculation of Aggregate Loss Distributions from Claim Severity Distributions and Claim Count Distributions," *PCAS LXX*, 1983. http://www.casact.org/pubs/proceed/proceed83/83022.pdf

3. Glenn Meyers, "Estimating Between Line Correlations Generated by Parameter Uncertainty," *CAS Forum*, Summer 1999a. http://www.casact.org/pubs/forum/99sforum/99sf197.pdf

4. Glenn Meyers, "Discussion of Aggregation of Correlated Risk Portfolios by Shaun Wang," *PCAS LXXXVI*, 1999b. http://www.casact.org/pubs/proceed/proceed99/99705.pdf

5. Glenn Meyers, "Setting Capital Requirements with Coherent Measures of Risk – Part 1 and Part 2," *Actuarial Review*, August and November, 2002. http://www.casact.org/pubs/actrev/aug02/latest.htm, http://www.casact.org/pubs/actrev/nov02/latest.htm.

6. Glenn Meyers, "The Economics of Capital Allocation," Presented at the Bowles Symposium, April 2003.

7. Stephen J. Mildenhall, "A Note on the Myers and Read Capital Allocation Method," submitted for publication, 2002. http://www.mynl.com/pptp/mr2.pdf

8. Stewart C. Myers and James A. Read, "Capital Allocation for Insurance Companies," *Journal of Risk and Insurance*, December 2001. http://www.aib.org/RPP/Myers-Read.pdf

9. Gary Venter, "Tails of Copulas," *Proceeding of the Casualty Actuarial Society LXXXIX*, 2002. http://www.cas.org/library/02pcas/venter.pdf)

10. Shaun S. Wang, Virginia R. Young, and Harry H. Panjer, "Axiomatic Characterization of Insurance Prices," *Insurance Mathematics and Economics 21* (1997) 173-182.

Appendix: Estimation of Correlation

Certainly one major driver of actuarial interest in correlation is the perception that positive correlation among lines of business, books of business, etc. has the potential to increase required capital. As a consequence of this observation, it seems to us that the program should be as follows:

- Estimate expected losses or loss ratios,
- measure deviations of the actuals from these expectations,
- and estimate correlations among these deviations as the correlations relevant to the required capital issue.

In an effort to parameterize various ISO models, we have carried out this program. For the sake of parsimony (to limit the required number of parameters to a relative few), we have imposed on correlation a model structure as documented in Meyers [1999a and 1999b]. We estimate correlations within company between lines of business and between company both within and between lines of business. These correlations among companies and among lines of business then drive correlations among reinsurance contracts written on those companies and lines of business.

Our dataset includes a fairly large number of companies, and our models are parsimonious in the sense of assuming that the same correlation model parameter values apply across all companies within a line of business. So our estimates are in effect pooled estimates. Even so, parameter estimates (contagions and covariance generators) were never more than two or three or four times their associated standard errors. Common statistical practice holds that an estimate is not statistically significant (at the approximately 95% level) unless the estimate in absolute value is at least twice its standard error. Had our dataset not included as many companies or had we attempted to estimate separate parameters by company (or at least by class of company), standard errors would have been larger in relation to their estimates. So it is doubtful that we would have found many parameter estimates significant at the 95% level. The large number of companies and the pooling are necessary to achieve significance.

The next section of this Appendix will address some philosophical issues of just precisely what correlation do we wish to measure anyway, and what are some of the adjustments we must make to observe this correlation. The following section will then discuss the correlation model of Meyers [1999a and 1999b] and an introduction to how we estimate the parameters appearing in the resulting formulae. The remaining sections will discuss the technical details of the estimation, with a few representative results presented at the end. We defer to the end of the model discussion a quick summary of the remainder of this Appendix, because even a quick summary of the technical details requires as background the topics we will discuss in the next two sections.

Correlation of What?

Suppose a realistic forecast, taking into account current rates and prices, estimates of trend, perceptions of current market conditions, etc., indicates that next year's losses will be higher than the long-term average. On the day the business is written, the insurance executive therefore already expects losses higher than average and makes some provision for that. Where the requirement for capital comes from, however, is the recognition that losses could emerge even higher than the already higher expected, and potentially higher than expected simultaneously for a number of lines of business, books of business, etc., due to positive correlation among those books. Thinking in this way clearly identifies the fallacy of measuring correlations of deviations about long-term averages, where some of the deviation from-long term average is due to predictable cycles, trends, etc. What matters, at least for correlation studies relevant to required capital, is not predictable deviation from long-term average but correlated, unpredictable deviation from expectations varying predictably from long-term averages.

As an enlightening thought experiment, consider an optimistic insurance company that consistently forecasts losses lower than their true expected value. Considerably more often than not, deviations of actual from forecast will be positive, yielding apparently fairly significant positive correlations among the outcomes, probably more positive correlation than would result if we were to measure deviations about true expected values. This thought experiment warns us that, to some extent, the correlations we measure will be dependent upon the way we estimate expectations from which we measure deviations.

As a further enlightening thought experiment, we ask what algorithm would most likely produce correlation estimates most relevant to the required capital issue. This would be the algorithm that most closely mimics the actual emergence over time of information in the insurance industry. Suppose for a number of companies and lines of business that we had time series of annual ultimate loss results (or results to date developed to ultimate), as well as potential predictor time series, such as losses emerged at each point in time (not developed to ultimate), rate and price indices, trend estimates at various points in time (based only on information up through that time), indicators of market competitiveness at various points in time, etc. As an example, suppose we sit at the end of year 10 and forecast year 11 based only on what the industry would have known at the end of year 10. Then in year 11 we calculate deviations of ultimate losses from these forecasts. Then we roll the time series forward to the end of year 11 and repeat the process, forecasting year 12, etc. Finally estimate correlations among these deviations.

The problems with this algorithm are at least twofold: 1) We probably need time series with duration of at least a couple of decades--at least the first decade to calibrate the time series forecasting model, plus at least another decade of forecasts from the calibrated model, and their attendant deviations and correlations, so that correlation estimates are not driven too much by events in any one year. In fact, it would probably be useful to have at least a couple of decades of forecasts and deviations so that we could potentially test the stability of correlation estimates over time. 2) We would need to reconstruct time series of what the industry knew at past points in time, such as rate and price indices, past estimates of trend, market competitiveness indices, etc. We might not be able to construct such time series at reasonable cost. Also, we might not be able to reconstruct other time series of what the industry knew or could have known at past times with any reasonable accuracy.

In light of these difficulties, we have constructed "forecasts" about which to measure deviations and correlations via an alternative algorithm. By line of business (LOB) and company, we have about a decade's worth of paid losses developed out to the oldest age in our available loss development triangles. We have not constructed time series of other potential predictors of those loss ratios. Instead, separately by LOB, we have developed generalized additive models for these loss ratios with main effects for company and a non-parametric, non-linear smoother term for year. The year effect is a loess smoother (Not a typo. Loess is a form of localized regression.) of local second degree with smoothing window over years sufficiently wide that long-term trends and turning points are captured without responding much to the random ups and downs of individual years. We have chosen a smoother of local second degree rather than first degree to better respond to turning points in the data.

The downside of this algorithm for correlation analysis is that the use of smoothers produces "forecasts" that, at any given point in time, depend on all of past, present, and future with respect to that point in time. Such "forecasts" may perform better than even the best of forecasts that must depend strictly on only the past, especially with respect to turning points and points of inflection. Therefore, some of what is captured in a smoother-based "forecast" (and therefore considered "predictable" with respect to that forecast) would be unpredictable and not captured by forecasts dependent strictly on the past and would instead be captured in the unpredictable deviations about those forecasts. Therefore, deviations about true forecasts dependent only on the past might tend to be somewhat larger and somewhat more correlated than deviations about smoother-based forecasts. As a consequence, our correlation estimates should be regarded as lower bounds.

On the other hand, the performance of our smoother-based forecasts may not be vastly superior to forecasts based only on the past that take advantage of more information than just losses, such as rate, price, trend, market competitiveness, etc. We would therefore not expect our correlation estimates to be vastly understated. Furthermore, we would expect those correlation estimates to be considerably closer to the mark than estimates based on deviations about long-term averages to the extent that in many of the lines we have studied there has been considerable long-term trend over the last decade; and we would argue that much of this long-term trend was indeed predictable, at least on a rolling one-year-ahead forecast basis.

A Correlation Model Based on Parameter Uncertainty

The reader is referred to Meyers [1999a and 1999b] where one of us has developed a model with correlation driven by parameter uncertainty. The essence of this model is captured in Simulation Algorithm #3 in the main text of this paper.

Losses are assumed conditionally independent; but correlation is imposed via severity multipliers assumed common across all lines of business and via frequency multipliers assumed common across all losses within a line of business and at least perfectly correlated, if not identical, across all lines within a so-called "covariance group." This model imposes a certain structure on correlations that depend upon parameters that can be estimated.

Although the models published in Meyers [1999a and 1999b] include both severity and frequency multipliers, we have chosen to fit to a version of the model with just frequency multipliers and have estimated the additional contribution to correlation from severity effects not by fitting data but rather by appeal to our understanding of severity-trend uncertainty. All losses across all lines are assumed multiplied by a common severity multiplier. This multiplier is a random variable with expectation 1 and variance b. If we assume our uncertainty regarding severity-trend translates to an uncertainty regarding severity on the order of 3%, then this translates to a b of approximately $(.03)^2 \approx 0.001$. Although we fit to a model form excluding severity-parameter uncertainty, the data we fit probably includes a component of correlation due to severity uncertainty, because we have certainly made no adjustments to the data to remove this particular uncertainty. Therefore, it is likely that the frequency uncertainty parameters of the model have taken up some of the slack and have responded to both frequency and severity uncertainty, at least to the extent that severity uncertainty can be captured by this model form. Then adding on top of frequency parameters, which may already have captured a portion of the severity effect, a b value estimated from first principles has the potential to overstate the total correlation. This is countervailing to the effect discussed in the previous section of this Appendix, which would potentially cause an understatement of correlation.

We note lastly that we have not yet studied correlations across years. But, within year, we note that we have studied across company/across LOB, across company/within LOB, within company/across LOB, and within company/within LOB (this last would be just variance, the usual process variance but augmented for the additional impact of parameter uncertainty).

Let L_{ijk} be the annual aggregate ultimate loss for line of business i, company j, and year k. Similarly for $L_{i'j'k}$. The two companies j and j' could be the same or different, the two lines i and i' the same or different. Assuming no severity parameter uncertainty, so $b = 0$, the covariance between L_{ijk} and $L_{i'j'k}$ is as given in Meyers [1999a]:

$$Cov[L_{ijk}, L_{i'j'k}] = \delta_{ii'}\delta_{jj'}\left\{\left(\frac{\sigma_i^2}{\mu_i} + \mu_i\right)E_{ijk} + (1+g_i)c_i E_{ijk}^2\right\} + \delta_{G_iG_{i'}}\sqrt{g_i g_{i'}}\,E_{ijk}E_{i'j'k}\,. \quad (A.1)$$

- $\delta_{ii'}$ is 1 if and only if $i = i'$ (i.e., the first and second LOBs are the same) and 0 otherwise. Likewise for $\delta_{jj'}$. In other words, the first term is nonzero only when first and second LOBs match, first and second companies match, and first and second years match, in other words, only when calculating variances.

- $\delta_{GiGi'}$ is 1 if and only if the first and second lines of business are in the same covariance group, otherwise 0. To get 1, first and second companies don't have to match, nor do first and second lines of business, but first and second lines of business have to be at least in the same covariance group.

- μ_i and σ_i are the mean and standard deviation of the severity distribution associated with LOB i.

- λ_{ijk} is the expected claim count associated with L_{ijk} and c_i is the contagion for LOB i, so the variance of claim count associated with L_{ijk} is λ_{ijk} + $c_i\lambda_{ijk}^2$.

- $E_{ijk} = E[L_{ijk}] = \lambda_{ijk}\mu_i$.

- g_i is the covariance generator associated with LOB i. In other words, in this line of business, parameter uncertainty associated with frequency is captured by a common multiplier across all companies within this line of business, the multiplier being a random variable with mean 1 and variance g_i. The formula above reflects one departure from the referenced Meyers [1999a and 1999b] papers. Whereas those papers assumed the same multiplier across all lines of business within covariance group, it is now assumed that across lines of business within covariance group the frequency multipliers could be different, with different covariance generators, but they are still assumed perfectly correlated. This results in replacing some occurrences of g_i in the earlier formulae with the $\sqrt{g_i g_{i'}}$ appearing above.

Recall that, by definition:

$$Cov[L_{ijk}, L_{i'j'k}] = E[(L_{ijk} - E[L_{ijk}])(L_{i'j'k} - E[L_{i'j'k}])].$$

Define the normalized deviation

$$\Delta_{ijk} = \frac{L_{ijk} - E[L_{ijk}]}{E[L_{ijk}]}.$$

Then divide through equation A.1 above by $E_{ijk}E_{i'j'k}$ to find:

$$E[\Delta_{ijk}\Delta_{i'j'k}] = \frac{\delta_{ii'}\delta_{jj'}\left(\frac{\sigma_i^2}{\mu_i}+\mu_i\right)}{E_{ijk}} + \delta_{ii'}\delta_{jj'}(1+g_i)c_i + \delta_{GiGi'}\sqrt{g_i g_{i'}}\,. \tag{A.2}$$

So, if $i = i'$ and $j = j'$, we are looking at a variance. Then that variance is a regression on $1/E$, with regression coefficient depending only on the parameters of the underlying severity distribution and with intercept term equal to $c_i + g_i + c_i g_i$. This term is approximately $c_i + g_i$ because the product $c_i g_i$ can be expected to be much smaller than either c_i or g_i, both of which are expected themselves to be small. If first and second companies are different but first and second lines of business are the same, then the expectation above is g_i, the covariance generator for the single common line of business. Regardless of whether first and second companies are the same or different, if first and second lines of business are different, then the expectation above becomes $\sqrt{g_i g_{i'}}$, the geometric average of the covariance generators of the two lines of business. If the two lines of business are in different covariance groups, then the expectation above is zero.

Suppose we estimate those expectations, and hence the parameters of our correlation model, from (weighted) averages of or regressions on pairwise products of normalized deviations of our underlying data. We will discuss the appropriate weights later. Consider first all pairwise products of normalized deviations where the first and second LOBs are equal to a single selected LOB of interest, with first and second companies different. From equation A.2, we expect an appropriately weighted average (across all companies and years) of these pairwise products to approximate the expectation g_i. We estimate $g_{i'}$ for a second LOB i' the same way. Having determined g_i and $g_{i'}$, suppose now we consider all pairwise products where the first LOB is i and the second is i', without constraint on first and second companies being the same or different. We expect that the appropriate weighted average of those pairwise products will be $\sqrt{g_i g_{i'}}$. If we find this indeed to be the case, then we conclude LOBs i and i' are in the same covariance group. But if we find the weighted average to be statistically insignificantly different from zero, we conclude that LOBs i and i' are in different covariance groups. Lastly, we consider pairwise products where the first and second company is the same and where the first and second LOB is the same and equal to a selected LOB of interest. According to equation A.2, these products should display a $1/E$ dependence. Regress these products on $1/E$ and identify the intercept estimate with $c_i + g_i$. Note that c never appears naked in these expressions, always in conjunction with g, but, having already inferred g, we can back out g to infer c.

For the rest of this Appendix we will carry out the following program:

1) In the next section, "Model for Expected Losses," we will discuss the estimation of the E_{ijk} and calculation of the normalized deviations Δ_{ijk} with an adjustment for degrees of freedom. The need for weights and the appropriate weights to use in modeling E_{ijk} will be important issues.

2) The following section, "Model for Loss Variances," will discuss the use of squared normalized deviations Δ_{ijk}^2 to fit the $1/E$ variance models mentioned above and estimate the sums of contagions and covariance generators by LOB, c_i + g_i.

3) The following section, "Other Pairwise Products," will discuss the use of other pairwise deviation products $\Delta_{ijk}\Delta_{i'j'k}$ with at least one of $i \neq i'$ or $j \neq j'$. Products in which the first and second LOBs are the same, $i = i'$, but companies are different, $j \neq j'$, yield estimators for the covariance generators g_i. Products in which the first and second LOBs are different, $i \neq i'$, provide

a test of whether two LOBs are in the same covariance group or not. The issue of weights will again be important. Also to be introduced at this point will be the use of the bootstrap to quantify standard errors of estimates.

4) The last section, "Some Representative Results," will discuss for two lines of business some representative results for contagion c_i, covariance generator g_i, and whether or not these two lines are in the same covariance group. Furthermore, for one of our representative lines, we will also perform the calculations measuring deviations relative to means not adjusted for long-term trends. We will indeed find much larger contagions and covariance generators. But, as we have already argued, these larger parameters are not appropriate for capital requirement calculations.

Model for Expected Losses

As already noted, we start with paid losses by LOB, by company (or company group), by year developed not to true "ultimate" but rather to the greatest age in loss development triangles available to us. We ratio these losses to premiums, build models for expected loss ratio, then multiply by premium to get back to estimates for expected loss. For each LOB, we actually test a number of denominators (premium, PPR, one or more exposure bases) in search of a denominator that produces a model for the ratio of loss to that denominator with a relatively high R^2. Presumably, for those denominators producing ratio models with lower R^2, the additional unexplained volatility is attributable to the denominator and interferes with good estimates for expected loss. High R^2 means the denominator is either stable or changes smoothly over time and is less likely to interfere with good estimates of expected loss.

Graph A.1.1 shows loss ratios by year, each line representing a separate company or company group. This is a package line with considerable property exposure, which may explain the apparent coordinated short-term up and down movement, which is evidence of correlation across company within LOB. The long-term apparent upward trend is probably just that, trend, was probably predictable, and, according to the discussion at the beginning of this Appendix, should not be considered evidence of correlation in the sense that we mean correlation.

Graph A.1.2 shows loss ratios by year for a liability line. Correlation is less readily apparent in this second graph. We should not be surprised if the correlation parameters we estimate for the second LOB are less than those for the first.

The graphs for these two lines are reasonably representative of graphs for the other lines we studied as well. The reader should note an important feature of these graphs that motivates the subsequent model. The lines for some companies lie consistently above the lines for other companies and appear to move in parallel to one another. Where correlation is visually significant (LOB 1), the parallel motion is evident even over short periods of time. Where correlation is less visually significant (LOB 2), the parallel motion is less pronounced over short periods of time but is still evident, on average, over the decade as a whole. This suggests a main-effects model with main effects for company and year. We assume no company/year interactions partly because such interactions are not apparent on the graphs and partly because we could argue that we lack sufficient data to estimate separate year effects by company anyway. We fit the year effect with a non-linear, non-parametric smoother to capture a wide range of possible behaviors across years – consistent trend, turning points, points of inflection, etc. This model produces fitted loss-ratio values that are parallel curves, a separate curve for each company.

The fitting is performed by invoking a generalized additive model package, specifying normally distributed errors, an identity link function, main effects for company and year, and a loess smoother on year with wide smoothing window (large "span"), so as not to respond too much to random hits in any one year. Although one could argue that, technically, loss ratios cannot be normally distributed (shouldn't be negative and are likely positively skewed), we observed deviations from normality sufficiently mild for our data that the normal assumption was acceptable, which brought us that much closer to the classic linear model. Also, we saw no evidence that the loss ratios themselves were not additive in the explanatory variables (company and year), hence the identity link function, which again brings us that much closer to the classic linear model. In fact, the only reason for invoking the generalized additive model, rather than the classic linear model, was our desire to impose a non-linear, non-parametric smoother on the year effect.

The generalized additive model was weighted. Over the years, it has been our experience fitting statistical models to insurance data that unweighted models are almost never appropriate. Weighted models are generally more appropriate, because insurance data points are almost never of equal credibility or volatility; and, furthermore, the range of credibilities or volatilities is sufficiently great that unweighted models are inadvisable. The general statistical practice is that the weight associated with a data point varies as the reciprocal of its variance. This practice produces minimum-variance fitted values. A general statistical rule of thumb is that, so long as the variances of the data points are sufficiently similar to one another (in other words, differ from one another by no more than a factor of two or three) and assuming the variances independent of the explanatory variables in the model, then the differences in results between a weighted and an unweighted model can be expected to be sufficiently modest that they are ignorable. Then an unweighted model is acceptable. The purpose of weighting is not to adjust for every last bit of difference in variance but rather to correct for gross asymmetries in variance. But most insurance data presents a range of variances considerably greater than a factor of two or three and so generally calls for the estimation of weighted models.

The classic actuarial assumption is that the variance of a loss ratio declines as one over some measure of volume, such as premium, which would suggest weighting on premium. But the formulas of the previous section of this Appendix would suggest that, in the presence of parameter uncertainty, the variance depends on two terms, one of form 1/*volume*, the second a constant greater than zero. So the very smallest risks, for which the 1/*volume* term dwarfs the constant, do indeed see a variance declining as 1/*volume*. The very largest risks, for which the 1/*volume* term has essentially died away to zero, see a variance essentially independent of size. If all the data is essentially small risks, weighting on volume is appropriate. If all the data is essentially large risks, doing an unweighted analysis is reasonable. Generally, we are somewhere in the middle, with risks all the way from the small to the large.

One possibility is to construct an iterated model. Select some weights. Fit a weighted model to find fitted means. Find the differences of actuals and fitted means, square the differences, and fit these squared differences to the variance model 1/*volume* plus a constant. Invert the fitted variances to find a new set of weights and iterate a few times. This is admittedly a fair amount of work. A "quick and dirty" alternative that we have frequently found to work adequately for weighting, where adequate means it removes gross asymmetries in variance without necessarily reducing all variances to exact equality, is to assume that variance dies away as 1 over some fractional power of volume; say, variance dies away as $1/\sqrt{volume}$--hence use the square roots of volumes as weights. Over quite a robust range of different models, we have found that this square root rule roughly captures the change in volatility from the small to the large.

As an example, Graph A.2 shows the same loss ratios as in Graph A.1.1 (LOB1), but plotted against premium rather than year. The smallest risks have premium as small as approximately $5 million. The largest premiums exceed $1 billion. So premium covers a range of two and a half orders of magnitude. As expected, loss ratio volatility appears to decline with increasing volume, but apparently not as fast as a 1/*volume* rule would imply. If the 1/*volume* rule held, as premium increased by more than a factor of 100, variances on the extreme right would be less than 1/100 of the variances on the extreme left, and standard deviations on the extreme right would be less than 1/10 of standard deviations on the extreme left. Standard deviations on the extreme left don't look 10 times as big as standard deviations on the extreme right--more like the three or four times as big that would be implied by variances that went as $1/\sqrt{volume}$; hence standard deviations that went as $1/\sqrt[4]{volume}$. So, in building our models for loss ratio for LOB 1, we have used weights of $\sqrt{premium}$. In other words, data points associated with the largest risks are assigned weights on the order of 10 times as large as data points associated with the smallest risks.

Graph A.3 shows the year effect for this model on LOB 1. The dotted lines are the fitted year effect plus and minus two standard errors, corresponding to an approximately 95% confidence interval. The year effect has been translated to yield an average effect of 0. The absolute level of loss ratios is captured by the other main effect, the company effect. So we see loss ratios have been trending upwards throughout the decade, increasing by more than 20 loss ratio points from the beginning to the end of the decade, but the trend has not been uniform throughout. There is a point of inflection at mid decade. Throughout the first half-decade, trend was positive but decreasing, until it vanished altogether at mid-decade, only to resume its upward movement at decade end. Because this happened to all companies (at least our model assumes so, being a main-effects-only model, but, as noted before, there is no evidence of different year effects by company), and because the trend was essentially consistently upward and of significant magnitude, if we were to measure deviations about the decade mean, we would find most deviations early in the decade negative, most late in the decade positive. We would infer considerably larger correlations from these deviations than from deviations measured about the varying-year effect plotted in Graph A.3. For illustrative purposes only, we have actually done both calculations and will report the results later in this Appendix.

This year effect has a cubic appearance. This shows the importance of the non-parametric component of the smoother on year. Because the smoother was locally quadratic, in the absence of a non-parametric component, the global year effect would have been linear or quadratic and could not have captured the pattern evidenced in Graph A.3. At the same time, the smoother is not so responsive as to pick up the year-to-year ups and downs apparent in Graph A.1.1. So long-term trends captured in the means, as driven by the year effect, therefore are removed from deviations about means, and don't impact correlation estimates. Short-term ups and downs are not captured in the year effect or the resulting means, so do flow through to deviations about those means and do carry through to correlations. This is the desired behavior.

Having identified good models for ratio of loss to one of premium, PPR, or exposure, we multiply the fitted values resulting from these models by the denominators to yield estimates for mean losses. These mean losses are then used to calculate the normalized deviations of the previous section of this appendix. As noted in the previous section, the normalized deviations are the actual loss minus the expected loss, the difference then divided by expected loss.

There is one additional, important adjustment to the normalized deviations not already discussed. These deviations are adjusted for degrees of freedom by multiplying by $\sqrt{n/(n-p)}$, where n, p, and the justification for this particular multiplier will now be described. Suppose the model for loss ratios for a particular LOB is based on n observed data points. The fitted model has p effective parameters, where p is the number of companies, plus two (because of the locally quadratic nature of the year smoother), plus the additional effective number of degrees of freedom of the non-parametric component of the year smoother, which was generally in the neighborhood of 0.8. An unbiased estimator for variance involves taking differences of actual and fitted values, squaring the differences, summing up the n squared differences, and dividing the sum not by n but by $n-p$. The way in which we subsequently use the normalized deviations to estimate correlation parameters amounts to taking averages, dividing sums of n terms by n rather than by $n-p$. By adjusting normalized deviations by the factor $\sqrt{n/(n-p)}$, we are adjusting squared deviations by $n/(n-p)$, the n's cancel, yielding the right denominator, $n-p$, in the end.

The need for applying a multiplier greater than 1 to the unadjusted normalized deviations can also be seen from the following argument, although this argument doesn't also establish the magnitude of the multiplier. We start with n data points.

To these data points we fit a model with p effective degrees of freedom. The fitted values are themselves random variables that approximate the "true" expected values to the extent that the model is the "true" model. But note that fitted values are pulled in the direction of the observed data and away from the true expected values by the fitting process (least squares, maximum likelihood, whatever). The magnitude of differences between actual and fitted values will therefore be smaller on average than the magnitude of differences between actual and true expected values. This shrinkage can be offset by multiplying the first differences by $\sqrt{n/(n-p)}$, where the actual value of the multiplier is established by the requirement that sums of squares reproduce the right unbiased estimate for the variance.

In the interests of wrapping up loose ends, we should note that, although we always started with a model with main effects for company and for year, with a smoother for year, the finally accepted models were many different variants on this. We sometimes found that company was not statistically significant; in other words, there was no statistically significant evidence that loss ratio differed by company. We sometimes found that the non-parametric component of the year effect was not significant, so the year effect was globally quadratic. Sometimes the quadratic term was not significant, so the year effect was globally linear (long-term constant trend). And sometimes even the linear effect was not significant, so there was no statistically significant evidence of loss ratio varying across years at all.

Model for Loss Variances

So now we have normalized deviations, adjusted for degrees of freedom. We consider all manner of pairwise products of these deviations. We demand that the year associated with the first factor in the pair match the year associated with the second factor, because we have not yet studied correlations across year. If we consider just those pairwise products where the first and second company also match, and where first and second LOB also match and are equal to some specified LOB of interest, then we are looking at squared deviations. Equation A.2 suggests that, if we plot these squared deviations against expected loss E, we should see a $1/E$ dependence plus a constant term, where the constant is the contagion plus the covariance generator for that LOB. See Graph A.4 for the graph just described for LOB 1. The circles represent the squared deviations from data. The triangles are the fitted values of the functional form $1/E$ plus constant.

The fit was created by least squares regression. There is again an issue of weights. Squared deviations for small expected loss appear considerably more volatile than squared deviations for large expected loss, and so should receive less weight. Otherwise, there is a considerable risk that some noisy data at small E could have a considerable impact on the estimate of the constant term out at large E. What weights might be appropriate? If the deviation Δ were approximately normal with standard deviation σ, then Δ^2/σ^2 would be distributed approximately chi-squared with one degree of freedom. This result would imply that Δ^2 has an expectation of σ^2 and a variance of $2\sigma^4$. In other words, the standard deviations of the squared-deviation random variables appear proportional to their expected values, which is not inconsistent with Graph A.4. This suggests the following algorithm. Fit the $1/E$ plus constant functional form to the squared deviations. Square the fitted values, take their reciprocals, and use these values as weights in another fit of the functional form to the squared deviations. Iterate a few times.

Other Pairwise Products

Consider next pairwise products where first and second year are the same, first and second LOB are the same and equal to some specified LOB of interest, but first and second company are different. These products measure correlation among companies within LOB, and their (weighted) average yields an estimator for the covariance generator for that LOB, per equation A.2. Consider first a plot of the second factor in each pair against the first factor in each pair. Can one visually see the correlation? Graph A.5.1 is such a plot for LOB 1.

The most striking thing about this plot is that the data appears to array itself in rows and columns. Consider an example. Suppose for this LOB we have 10 years, 10 companies, hence 100 independent observations from which we construct 100 normalized deviations. For each of the 100 deviations thought of as the first factor, there are nine deviations available as second factor (same year, each of the other nine companies), hence a total of 900 pairwise products relevant to this section of the Appendix (same year, different companies) and 900 plotted points on the plot of second factor vs. first factor of the form of Graph A.5.1. The points in this plot array themselves in columns of nine points and rows of nine points. The columns of nine result because all nine share the same first factor (plotted on the x axis) while the second factor (plotted on the y axis) ranges over nine possible values. Rows of nine also result because all nine share the same second factor while the first factor ranges over nine possible values. The nine points in a column are not independent but highly interdependent through their shared first factor. Likewise, the nine points in a row are not independent but highly interdependent through their shared second factor. These interdependencies through shared first and second factors apply also to the 900 pairwise products. It would be very wrong to treat these 900 pairwise products as 900 independent draws from some underlying process. This observation will be relevant to a later discussion of standard errors of parameter estimates, such as estimates of covariance generators.

Returning to Graph A.5.1, note the slightly tilted horizontal line. This is an unweighted linear regression line on the plotted points. It is included as an aid to visualizing a possible tilt to the plot, which would be indicative of a correlation, but the degree of tilt of this regression line is not a good estimator of the correlation. First, points with either very low or very high first deviation may be highly leveraged and highly influential in estimating the unweighted regression line. Yet these extreme first deviations are likely to be the most volatile and the least deserving of receiving any significant weight. An unweighted regression gives them too much weight. Second, the regression line treats all the plotted points as independent of one another, and we have already argued that there is a great deal of interdependency among these points. So the plotted regression line should be treated as a visual aid only and not considered a good estimator. We have argued in a previous section of this Appendix that a weighted average of pairwise products, with judicious choice of weights, might be a good estimator of covariance generators.

The deviations of Graph A.5.1 are those measured about expected losses taking into account the year effect of Graph A.3. As an additional aside on the potential distortion of estimating correlations from deviations about grand means, Graph A.5.2 shows a plot corresponding to Graph A.5.1 of deviations vs. deviations, measured about expectations not reflective of the year effect. The apparent correlation is much greater, the excess correlation being driven by the failure to remove long term predictable trend from the deviations.

We have concluded that, because of various technical difficulties, plots of deviations vs. deviations of the form of Graph A.5 are useful visual aids but not good estimators. As weighted averages of pairwise products of deviations can be used as estimators, what weights are appropriate? Previously, we presented a heuristic argument in terms of the chi-squared distribution for squared deviations; in other words, for pairwise products where the first and second factors are identical. But we don't know what the sampling distribution might be for pairwise products of deviations where the first and second factors may be interdependent but not identical. Suppose we plot pairwise products against some measure of volume to see if there is any evidence of changing volatility with increasing volume. For each of the first and second factors of a pairwise product, there is a measure of volume, namely the expected loss associated with that deviation, but the two expected losses are unlikely to be equal. Suppose we define as a measure of volume for the pairwise product the geometric average of the expected losses for the first and second deviations in the product; in other words, the square root of the product of the two expected losses. Call it E.

Graph A.6.1 shows a plot for LOB 1 of the pairwise deviation products, same year first and second factors, different companies, against this volume measure E. Pairwise products associated with larger volumes are clearly less volatile and so should receive more weight in any weighted average of these products. Suppose we imagine that the variance of the sampling distribution of a pairwise product declines as unity over some power of E. Dividing the observed pairwise products by the square root of the presumed variance law and plotting this against E should produce a graph more symmetrical left to right than Graph A.6.1. Suppose we guess the variance law to be $1/E$. Then multiply pairwise products by $\sqrt{E}$. Graph A.6.2 shows this plot. We have gone from a graph that shows more volatility on the left to one that shows more volatility on the right. Clearly, a $1/E$ variance law overdoes it. Suppose we assume a variance law $1/\sqrt{E}$. Then multiply pairwise products by the fourth root of E. Graph A.6.3 shows the resulting plot is far more symmetric than either A.6.1 or A.6.2, supporting a variance law something like $1/\sqrt{E}$ and, therefore, a weighted average of pairwise deviation products with weights proportional to $\sqrt{E}$ as a reasonably best estimator from among this family of estimators of the covariance generator for this LOB.

Now that we have an estimate for the covariance generator, how precise is it? What is the standard error of that estimate? Generally, when an estimator is a weighted average of independent observations, the standard error of the estimate is the standard deviation of one observation divided by the square root of the number of observations, with some adjustment for the weighting. As we have already argued, these pairwise products are far from independent of one another, ruling out the square root of n rule. We have chosen to estimate standard errors of estimators via bootstrap. From the original data draw a data resample of the same size as the original data set, but with replacement, so that some data points might not appear at all in the resample and others might appear more than once. Re-estimate the statistic or parameter of interest from this resample. Repeat this many times, building up a collection of estimates, from which collection one can estimate such quantities as the standard deviation and extreme percentiles of the estimator. Statistical rules of thumb suggest that, whereas one may need hundreds of resamples to reasonably estimate extreme percentiles (such as the 95th or 99th) of the sampling distribution of the estimator of interest, as few as fifty resamples will yield a reasonable estimate of the standard error of the estimator.

Furthermore, to preserve the two-way structure of the underlying problem on company and year, as well as to estimate the relative impact of company and year on estimators, we bootstrap separately on company and year. Bootstrapping on company yields a standard error of the estimator due to the randomness of which companies are in or out of the database. In other words, if certain companies were dropped from the database, and certain others were added, how much could we expect the estimator to vary from its current value? Bootstrapping on year yields a standard error of the estimator due to the randomness of which years are in or out of the database. The total standard error of the estimator is the square root of the sum of squared standard errors due to company and year separately.

An example may again be useful. Suppose our previous example with an LOB with ten years and ten companies. This produces 100 normalized deviations, 100 squared deviations used to estimate the variance model, and 900 pairwise deviation products, first and second years the same, first and second LOBs the same and equal to the LOB in question, but different first and second companies, from which an estimate for the LOB covariance generator is calculated. One way to bootstrap would be to draw from the 100 deviations with replacement, but it is likely that this would produce a resampled dataset in which some years were represented by some companies but not all ten companies, and some companies were represented by some years but not all ten years. The resampled dataset would not preserve the two-way structure of the original on company and year. Also, from this resample it would be impossible to segregate the potentially interesting different impacts of company and year.

We chose to resample on company and year separately. One resamples on company by drawing ten companies with replacement from the original list of ten. As an example, the resampled list might include eight of the original ten appearing once each, the ninth appearing twice, and the tenth not at all. Then one takes all ten years for each of the resampled companies. The result would be 100 deviations, the first 80 from the original 100 representing the first eight companies, then 81 through 90 from the original 100 representing the ninth company, then 91 through 100 repeating 81 through 90, representing the ninth company showing up a second time in that particular resampling on company. So, although the resample includes 100 deviations, there are only 90 distinct values, because company 9 occurs twice in the resample. One uses these resampled 100 deviations to calculate the previously discussed variance model and covariance generator estimator. Resample 50 times to estimate standard errors for the estimators.

Next resample on year by drawing ten years with replacement from the original list of ten. As an example, the resampled list might include six of the original ten appearing once each, the seventh and eighth appearing twice each, and the ninth and tenth appearing not at all. Then take all ten companies for each of the resampled years. The result would be 100 deviations but only 80 distinct values, because years 7 and 8 occur twice in the resample. Use these resampled 100 deviations to calculate the previously discussed variance model and covariance generator estimator. Resample 50 times to estimate standard errors for the estimators.

The previous section of this Appendix, on the variance model, considered pairwise deviation products where the first and second factor years were the same, first and second LOBs the same, and first and second companies the same; in other words, the pairwise products were actually squared deviations. These lead to variance models and estimators for the sum of contagion and covariance generator for the LOB. In this section, we have considered pairwise products with first and second years the same, first and second LOBs the same, but first and second companies different. These products lead to estimates of correlation among companies within LOB, to estimators for the LOB covariance generator. Other pairwise products not yet discussed but of potential interest would be those for which first and second years are the same, but first and second LOBs are different. Such products would lead to estimates of between-LOB-correlation, to estimators for the geometric average of the covariance generators for the two LOBs if they are in the same covariance group, or to a statistic not statistically different from zero if the LOBs are in different covariance groups. We will not discuss these products further other than to note that the weighting and bootstrap issues discussed above are the same for these products and were addressed in the same way.

Some Representative Results

Before discussing Exhibits A.1 through A.3, which provide some representative results, we should note that we tested two other model issues that have not yet been discussed.

1) Between company pairwise deviation products yield estimators for covariance generators. We asked whether there was any evidence that these covariance generators varied by size of company. We tested this by regressing the appropriate pairwise products against the base 10 logarithm of the size of the company, size measured as the expected loss for that LOB. A statistically significant regression coefficient for the log explanatory variable would have been evidence of a size dependency. A statistically significant positive coefficient would have been evidence of a covariance generator increasing with increasing company size, and vice versa for a statistically significant negative coefficient. We used log(size) as the explanatory variable on the assumption that the effect, if there was one, would be logarithmic in size, that the magnitude of the effect would be about the same when going from a company of size 1 to size 10 as when going from a

company of size 10 to one size 100, etc. No statistically significant size effects for the covariance generators were detected.

2) For certain property lines, we asked whether much of the apparent correlation arose through catastrophes. We eliminated the heavy catastrophe years of 1992 and 1994 and found that correlations did indeed go down but were still significant.

Turning now to Exhibit A.1, this exhibit considers just pairwise deviation products where first and second LOB are LOB 1. Considered first are products where first and second companies are different ("Between companies"), hence the expectation is g_1. Based on a weighted average of the relevant pairwise products from the data, the point estimate for g_1 is 0.0026. The square root of this, 0.051, is the standard deviation of the underlying frequency multiplier, which appears to indicate a frequency parameter uncertainty impacting LOB 1 industry wide of on the order of plus or minus 5%. Bootstrapping on years yields a range of estimates for g_1 with a standard deviation of 0.0008. Bootstrapping on companies yields a standard error due to companies of 0.0009. So uncertainty regarding this parameter due to years is comparable to the uncertainty arising through companies. The total standard error for g_1 is a combination of standard errors due to years and companies and is 0.0012. The estimate for g_1 is more than twice its standard error, so is certainly statistically significant.

The test for g_1 size dependence yields a regression coefficient for the *log*(*size*) explanatory variable of -0.00004, with a standard error estimated from bootstrap of 0.00344. The standard error is much larger than the parameter estimate. There is no statistically significant evidence that g_1 depends upon size.

Considering next pairwise products with first and second LOB equal to LOB 1 and with first and second companies equal ("within company"; in other words, the squared deviation products) yields an estimate for LOB 1 of contagion plus covariance generator of 0.0226 with a standard error of 0.0092. This is certainly significant. The difference of the $c + g$ estimate (0.0226) and the g estimate (0.0026) yields an estimate for the contagion c for LOB 1 of 0.0200.

If, just for the sake of illustration, not that we argue this is the right thing to do, we repeat these calculations for LOB 1 using deviations about grand means rather than about means adjusted for the year effects of Graph A.3, we find much larger correlation estimates. For g_1, instead of the 0.0026 with standard error 0.0012 discussed above, we find 0.0135 with standard error 0.0051. This latter value for g_1 implies a frequency parameter uncertainty of 11.6% vs. the 5% discussed above. Likewise, for $c_1 + g_1$, instead of the 0.0226 with standard error 0.0092 discussed above, we find 0.0298 with standard error 0.0099. Failing to adjust deviations for long-term predictable trends significantly inflates correlation estimates in ways not directly relevant to the required capital issue.

Exhibit A.2 shows the same statistics for LOB 2, a g estimate of 0.0007 with standard error of 0.0004 (hence just about significant at two standard errors, indicating a frequency parameter uncertainty of plus or minus 2.6%), no significant size dependence of this g estimate, and a significant estimate of $c + g$ of 0.0090 with standard error of 0.0023. From comparing Graphs A.1.1 and A.1.2 we had suspected we would find more correlation in LOB 1 than in 2, and indeed we find g for LOB 1 larger than that for LOB 2. $c + g$ measures large risk volatility (the limit as the $1/E$ term dies away). This is also larger for LOB 1 than for LOB 2.

Turning lastly to Exhibit A.3, this considers pairwise products where the first LOB is LOB 1 and the second LOB is LOB 2, hence measures between LOB correlations. This yields an estimate of $\sqrt{g_1 g_2}$ of 0.0005 with a standard error of 0.0006. Because this statistic is not statistically significantly different from 0, there is no evidence that LOBs 1 and 2 are in the same covariance group. Knowing what lines of business LOB 1 and 2 are, we did not expect them to be in the same covariance group and are not surprised by this result.

Loss Ratios by Company (LOB 1)

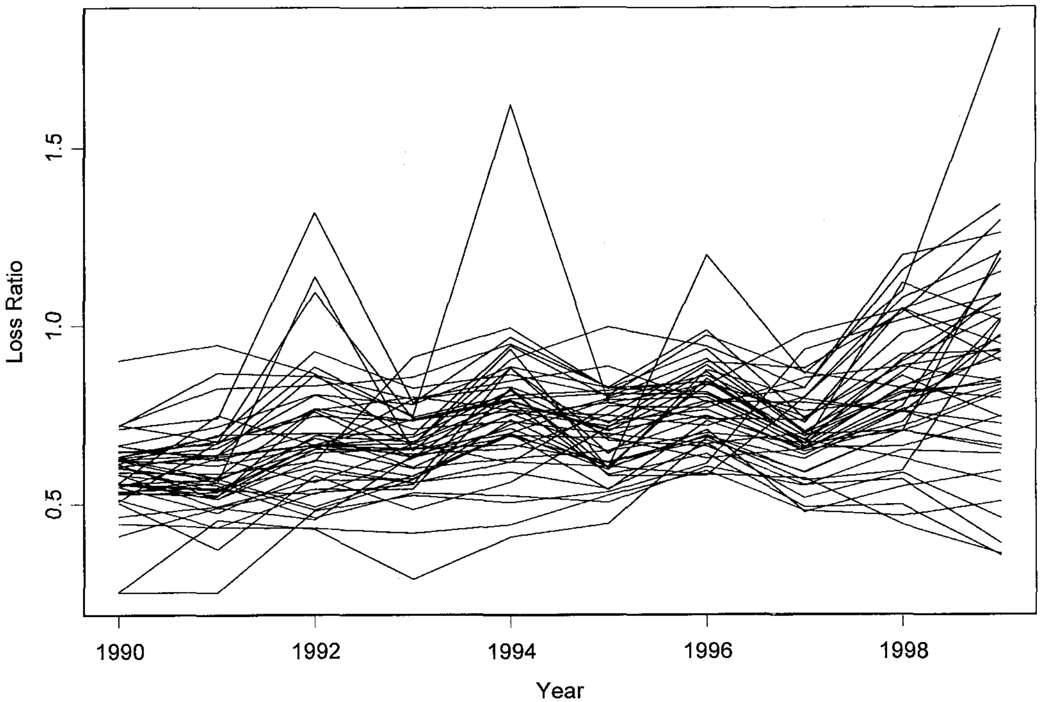

Graph A.1.1

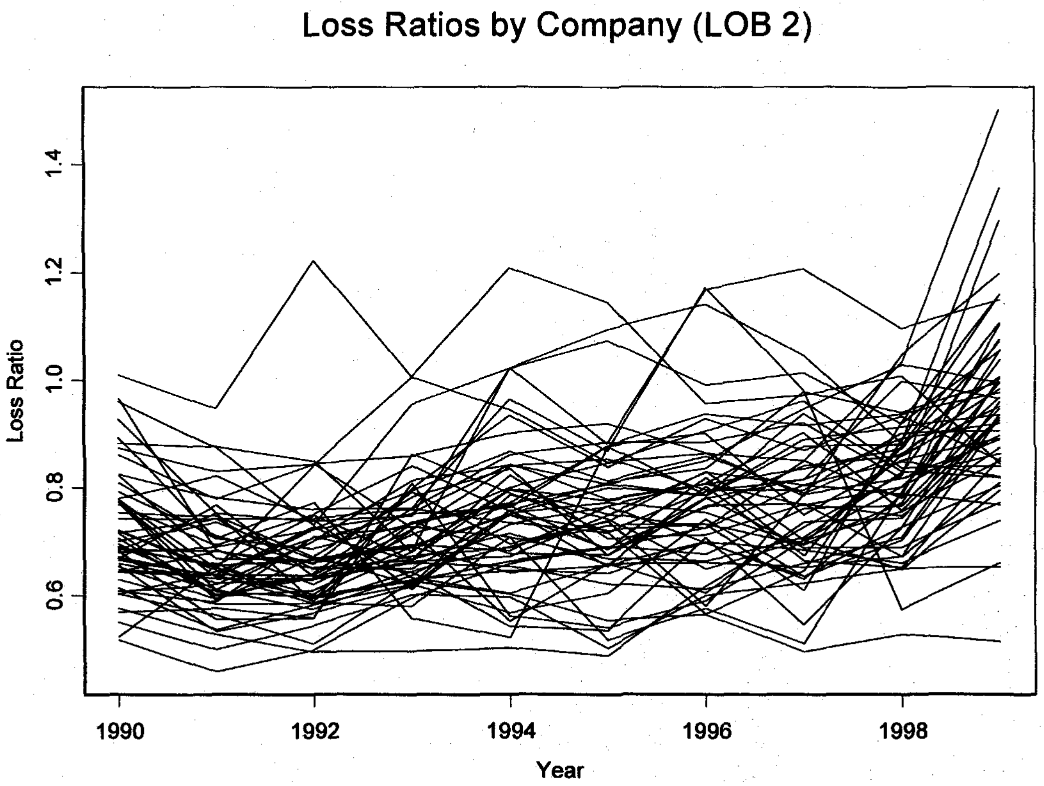
Loss Ratios by Company (LOB 2)
Loss Ratio
0.6
0.8
1.0
1.2
1.4
1990
1992
1994
1996
1998
Year

Graph A.1.2

Loss Ratio vs. Premium Volume (LOB 1)

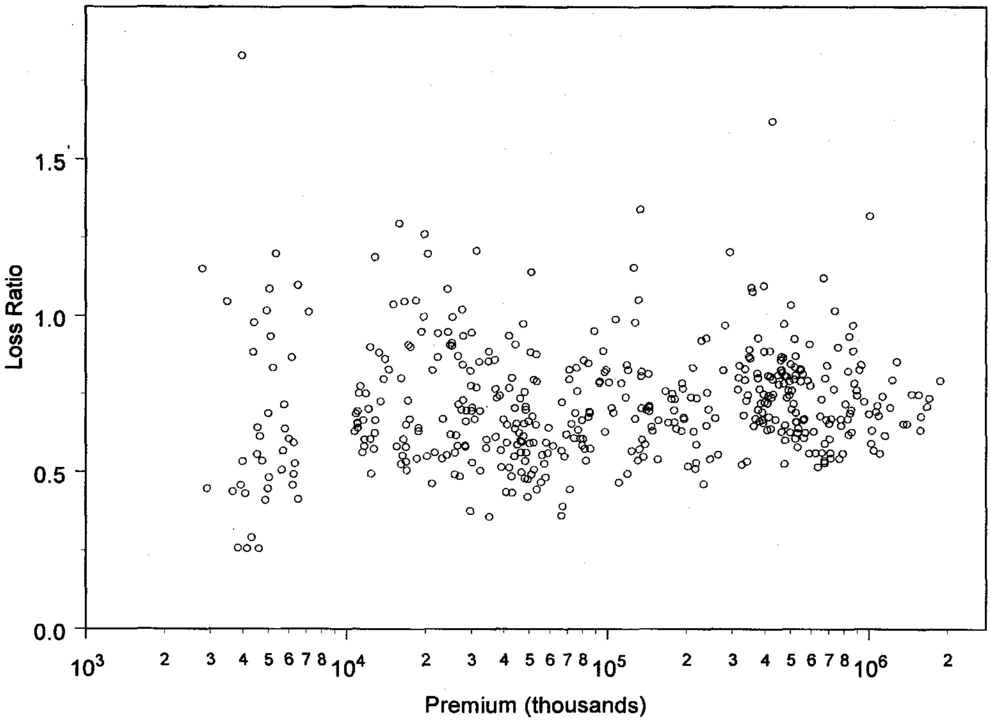

Graph A.2

Loss Ratio Year Effect (LOB 1)

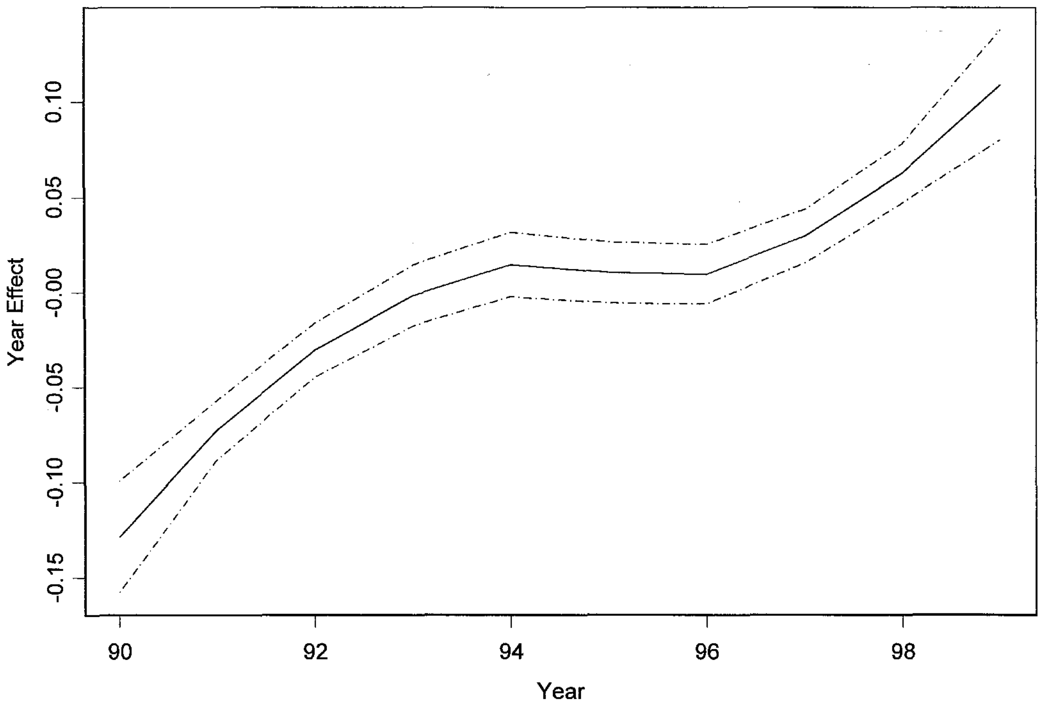

Graph A.3

Squared Deviation vs. Expected Loss (LOB 1)

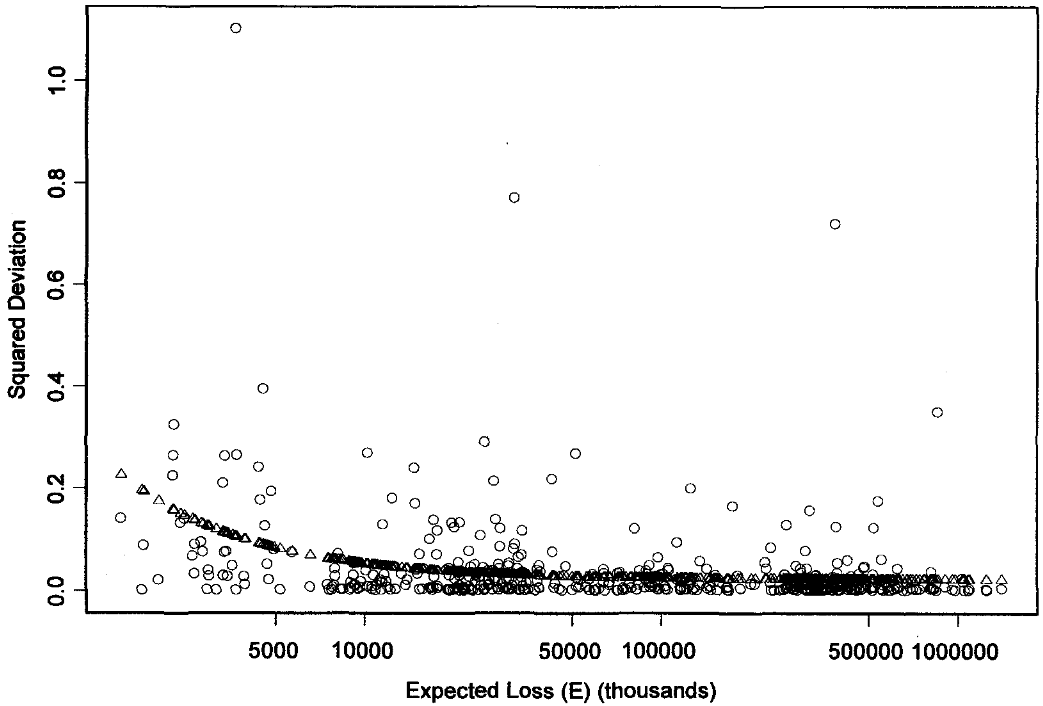

Graph A.4

Deviation vs. Deviation (LOB 1, Full Trend Model)

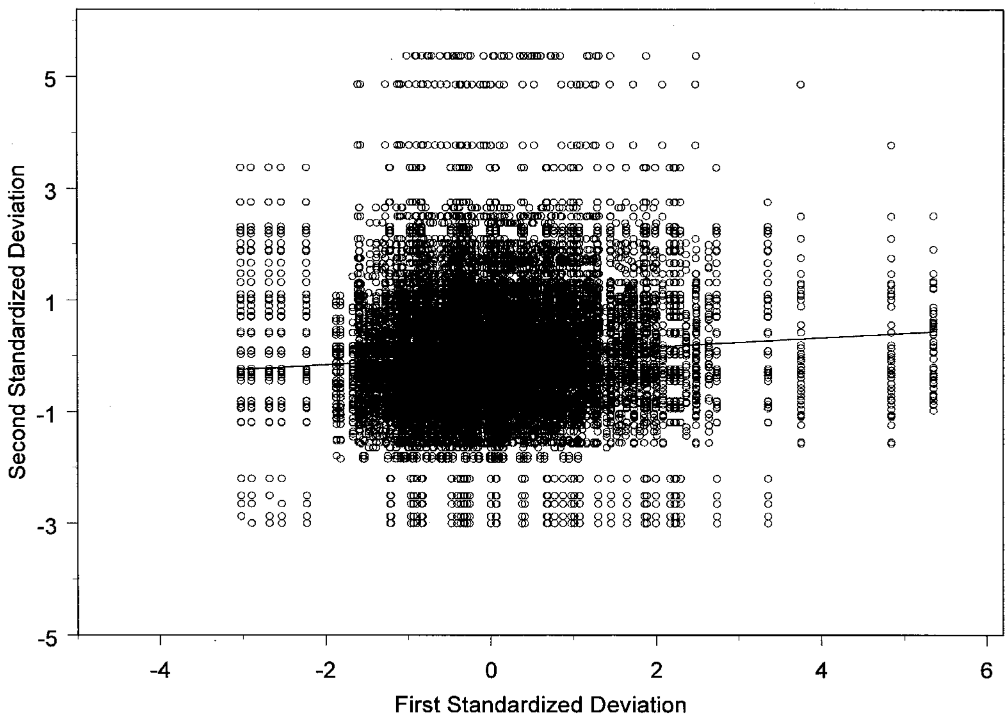

Graph A.5.1

Deviation vs. Deviation (LOB 1, No Trend Model)

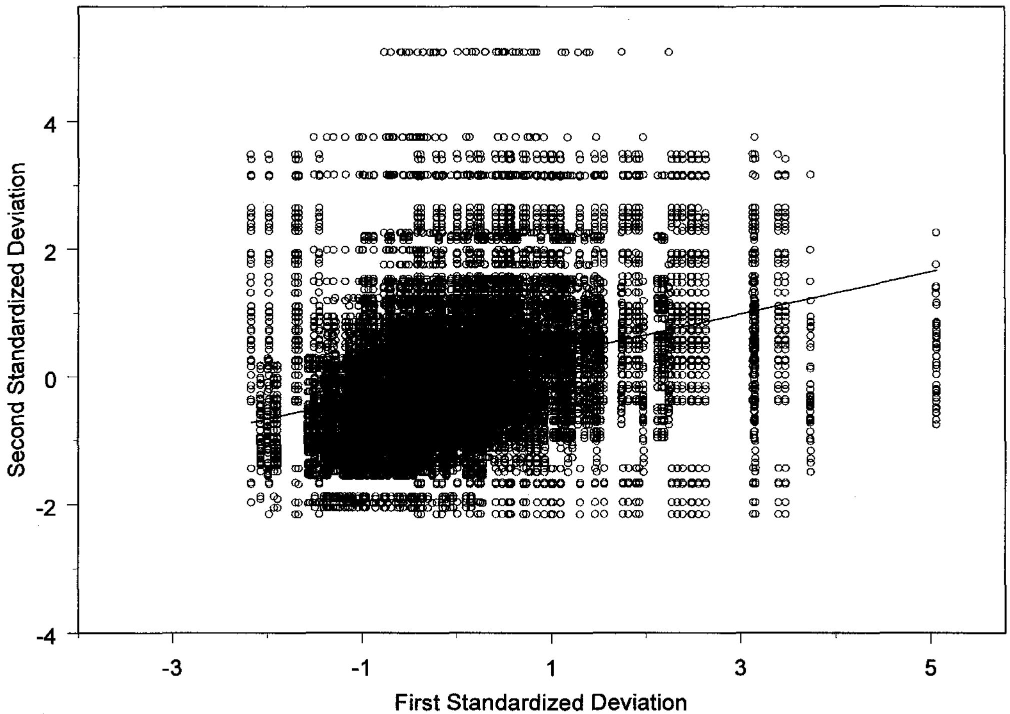

Graph A.5.2

Pairwise Deviation Products vs. E (LOB 1)

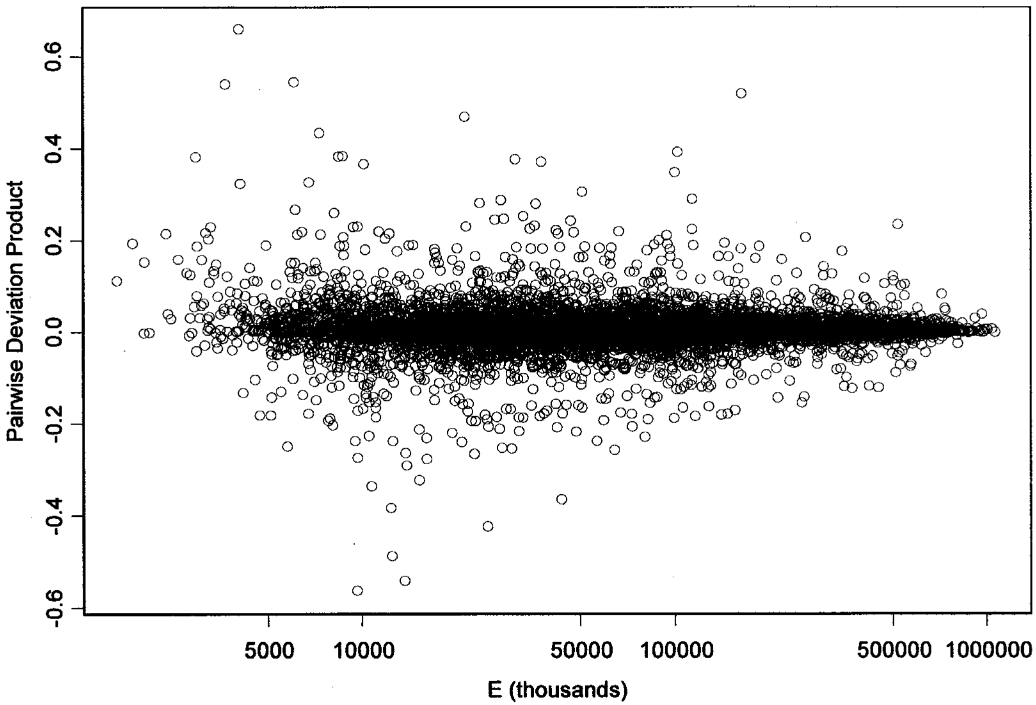

Graph A.6.1

Pairwise Deviation Products vs. E (LOB 1)

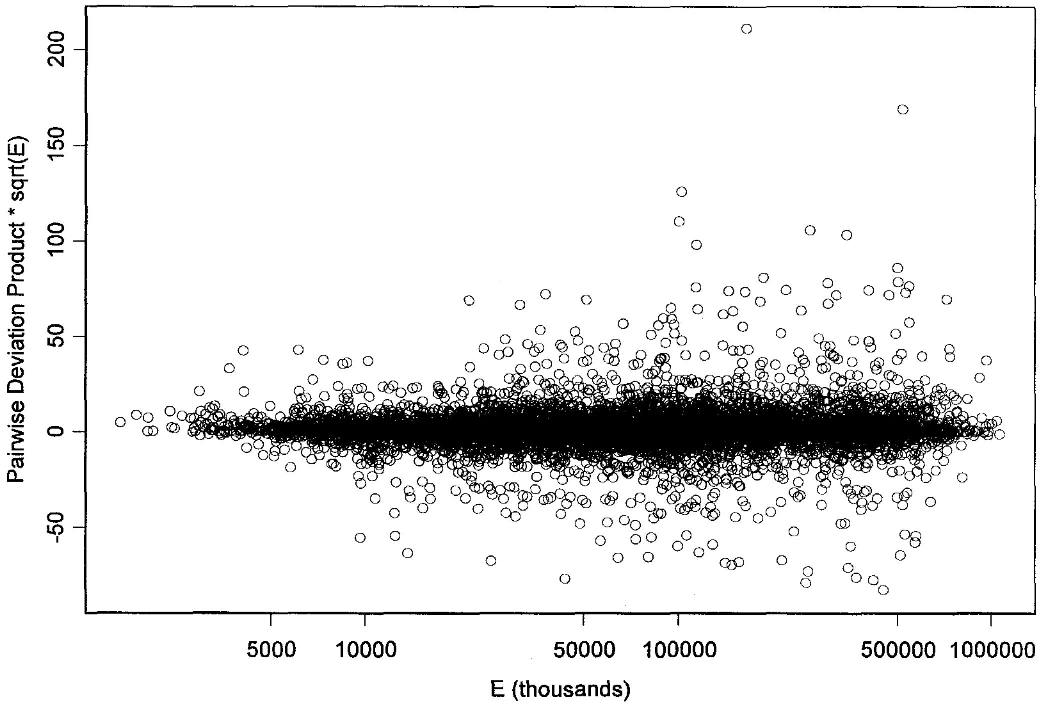

Graph A.6.2

Pairwise Deviation Products vs. E (LOB 1)

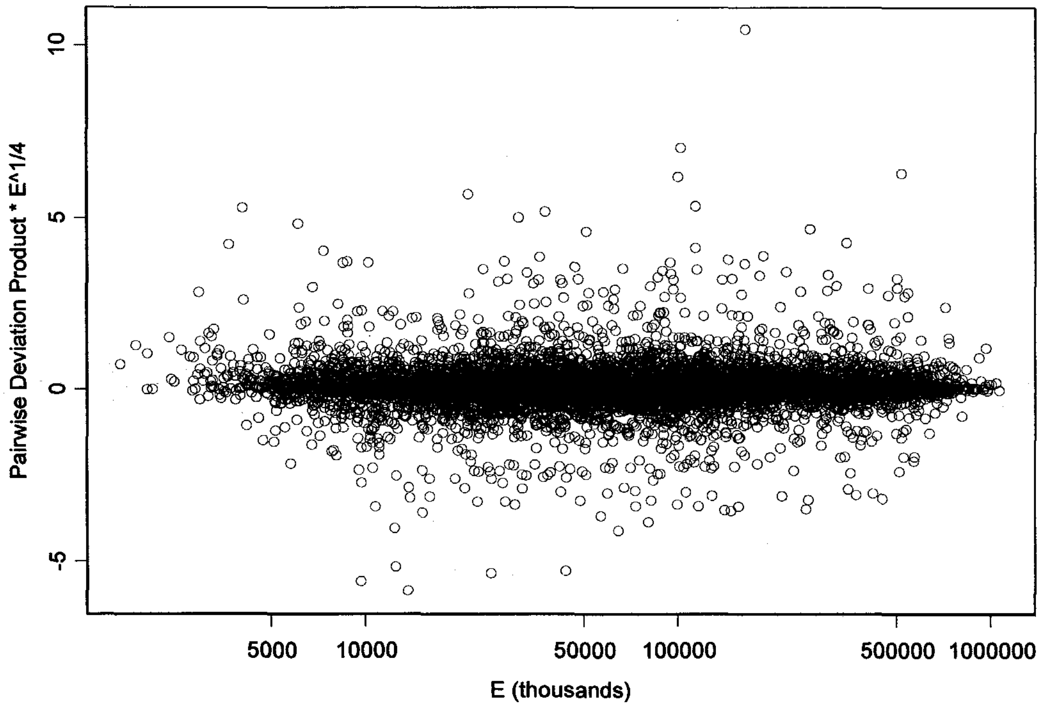

Graph A.6.3

Exhibit A.1

Correlation Parameter Estimates

LOB 1

Between companies: g

Estimate: 0.0026

Standard error due to years: 0.0008

Standard error due to companies: 0.0009

Full standard error: 0.0012

Between companies: $\log_{10}$(size) coefficient

Estimate: -4e-005

Standard error due to years: 0.00235

Standard error due to companies: 0.00251

Full standard error: 0.00344

Within company: $c + g$

Estimate: 0.0226

Standard error due to years: 0.0048

Standard error due to companies: 0.0078

Full standard error: 0.0092

Exhibit A.2

Correlation Parameter Estimates

LOB 2

Between companies: g

Estimate: 0.0007

Standard error due to years: 0.0002

Standard error due to companies: 0.0003

Full standard error: 0.0004

Between companies: $\log_{10}$(size) coefficient

Estimate: -0.00065

Standard error due to years: 0.00050

Standard error due to companies: 0.00065

Full standard error: 0.00082

Within company: $c + g$

Estimate: 0.0090

Standard error due to years: 0.0007

Standard error due to companies: 0.0022

Full standard error: 0.0023

Exhibit A.3

Correlation Parameter Estimates

LOB 1 vs. LOB 2

Between and within companies: *g*

Estimate: 0.0005

Standard error due to years: 0.0005

Standard error due to companies: 0.0003

Full standard error: 0.0006

Between and within companies: $\log_{10}$(size) coefficient

Estimate: -0.00086

Standard error due to years: 0.00080

Standard error due to companies: 0.00106

Full standard error: 0.00133