INSURANCE APPLICATIONS OF BIVARIATE DISTRIBUTIONS

DAVID L. HOMER AND DAVID R. CLARK

Abstract

A technique is demonstrated for aggregating bivariate claim size distributions using a two-dimensional Fast Fourier Transform. Three insurance applications are described in detail relating to: 1) individual risk rating, 2) loss and allocated expenses, and 3) Dynamic Financial Analysis.

ACKNOWLEDGMENTS

The authors wish to thank the following individuals: Dan Heyer, Hou-Wen Jeng, Don Mango, Jim Sandor, and Chuck Thayer.

1. INTRODUCTION

1.1. The Basic Problem

When pricing insurance contracts it is useful to estimate not only the average insured loss but also the insured loss distribution. Although an initial approach may include only an estimate of the mean, risk measures generally require an estimate of the distribution. This problem is often solved by modeling losses as a sum of individual claims. A frequency distribution describes the number of claims N; a severity distribution describes the size of each claim X_k. The individual claim sizes are usually assumed to be independent and identically distributed (iid) as well as independent from the claim counts. This model is known as the *Collective Risk Model* [3]. The aggregate loss dollars Z are the sum of the individual claim sizes

$$Z = X_1 + \cdots + X_N. \tag{1.1}$$

The expectation and variance of Z are easily expressed in terms of the frequency and severity components

$$\mathrm{E}(Z) = \mathrm{E}(X)\mathrm{E}(N). \tag{1.2}$$

$$\mathrm{Var}(Z) = \mathrm{Var}(X)\mathrm{E}(N) + \mathrm{E}(X)^2\mathrm{Var}(N). \tag{1.3}$$

Estimating the aggregate loss distribution requires more work, but there are numerous techniques available: simulation, Fast Fourier Transform, continuous Fourier Transform [1], recursion [4, 8], and moment matching [5, 9]. In this paper, the Fast Fourier Transform (FFT) will be used. The FFT has been described in detail by Robertson [7] and Wang [10], and an overview is also included here as Appendix C.

1.2. A Problem That Includes Dependencies between Loss Components

The collective risk model as outlined above is sufficient to describe most insurance policies. One example in which this model is not sufficient arises in individual risk rating. A policy may provide specific excess coverage above a per-occurrence retention, and may also provide coverage in excess of an aggregate amount for the retained losses. The excess of aggregate cover is commonly called a *stop loss* cover.

The distributions for either the specific excess or stop loss covers can be estimated using the collective risk model. However, it is more difficult to estimate the distribution for the sum of the two covers because there is a dependence between the pieces. One trivial element of the dependence is easily seen—if there are no retained losses then there are no losses in excess of the retention.

Section 2 provides a more detailed description of this problem.

1.3. Aggregating with the FFT—A Brief Review

Before introducing the complication of the dependence between two coverages, we will briefly review the Fast Fourier Transform (FFT) technique for evaluating a standard collective risk model. Appendix C provides a more detailed review.

In order to compute the aggregate loss distribution using the FFT, the severity distribution is expressed as a probability vector[1] $x = (x_0, x_1, \ldots, x_{n-1})$. Each element x_k is the probability of a claim having size ck, where c is a scaling constant.

The distribution of the claim counts N is incorporated with the use of its Probability Generating Function (PGF)

$$\mathrm{PGF}(t) = \mathrm{E}(t^N). \tag{1.4}$$

The frequency and severity components are put together using a standard FFT technique. Denoting the FFT and its inverse as FFT(x) and IFFT(x), respectively, the probability vector for the aggregate losses is computed as

$$z = (z_0, z_1, \ldots, z_{n-1}) = \mathrm{IFFT}(\mathrm{PGF}(\mathrm{FFT}(x))). \tag{1.5}$$

The PGF is applied *elementwise*, i.e., with some abuse of notation,

$$\mathrm{PGF}((t_0, t_1, \ldots, t_{n-1})) = (\mathrm{PGF}(t_0), \mathrm{PGF}(t_1), \ldots, \mathrm{PGF}(t_{n-1})). \tag{1.6}$$

The vector size n must be large enough that the probability of aggregate losses greater than cn is negligible. Any probability mass for losses greater than cn will wrap around, i.e., mass for losses greater than cn will be treated as though it is mass for the available claim sizes $(0, c, 2c, \ldots, nc)$. The wrap-around problem is typically avoided by padding the vector with zeros as discussed in Robertson [7] and Wang [10].

---

[1] x is indexed starting at zero. x_0 is the probability of a claim of size zero.

1.4. Building a Bivariate Loss Distribution

The goal is to obtain a bivariate distribution of aggregate retained losses and aggregate excess losses. This will be represented as a probability matrix[2] M_z where $M_z(j,k)$ is the probability that aggregate retained losses are $c_1 j$ and aggregate excess losses are $c_2 k$. As before, c_1 and c_2 are constant scale factors.

For a single claim this matrix is easily constructed. Suppose $x = (.4,.3,.3)$ and $c = 1{,}000$. Then for a 1,000 deductible, with $c_1 = c_2 = c = 1{,}000$,

$$M_x = \begin{bmatrix} .4 & 0 & 0 \\ .3 & .3 & 0 \\ 0 & 0 & 0 \end{bmatrix}. \tag{1.7}$$

The matrix M_x fully specifies the probabilities and dependencies of losses in the retained and excess layers. The sum across rows $(.4,.6,0)$ produces the distribution of the retained losses; the sum down the columns $(.7,.3,0)$ produces the distribution of the excess losses.

The advantage at this point is that the same FFT technique can be used to calculate aggregate losses for M_x that we used to calculate aggregate losses for x. With FFT() and IFFT() now representing the two-dimensional FFT and its inverse, and with PGF() as before, we compute the aggregate loss matrix M_z

$$M_z = \text{IFFT}(\text{PGF}(\text{FFT}(M_x))). \tag{1.8}$$

As in the one-dimensional treatment, the PGF is applied elementwise and the matrix M_x must have sufficient padding so that M_z can hold the significant mass. Appendix A provides an example of the two-dimensional FFT using publicly available software.

The FFT technique is not the only way to aggregate M_x. Sundt [8] shows that M_x can be aggregated using a recursive technique.

[2] M_z indices start from zero.

The aggregation of bivariate severity matrices can be applied to other problems as well. In what follows, three specific examples will be explored. In the first, the combined distribution of losses on specific excess and aggregate excess is considered. In the second, bivariate loss and ALAE distributions are computed, and in the third example, a problem with a simulation technique often used in DFA analysis is reviewed and corrected.

2. PER-OCCURRENCE AND EXCESS-OF-AGGREGATE COVERS IN INDIVIDUAL RISK RATING

The first problem that we will review is common in individual risk rating.

A fictional large insured, Dietrichson Drilling, is interested in retaining the majority of their "predictable" workers compensation losses, and mainly seeks to purchase insurance to cover individual large claims. For example, they may choose to retain the first 600,000 of each loss occurrence. At the same time, they may have a concern that the number of occurrences could also be higher than expected, and therefore seek protection on the total dollars of retained loss.

Our company, Pacific All Risk Insurance Company, has been asked to provide coverage on a per-occurrence basis of 400,000 excess of 600,000, and then also a stop loss cover to pay in the event that their total retained loss exceeds 3,000,000. The underwriter at Pacific All Risk has proposed the structure shown in Table 2.1.

As the Pacific All Risk actuary, you have selected frequency and severity distributions, and have estimated the expected losses for each of these coverages. In order to calculate the needed risk load on the program, however, you need to estimate the distribution of the sum of the two coverages.

The company's Fast Fourier Transform (FFT) model allows you to estimate a distribution for either the per-occurrence or the

TABLE 2.1

POLICY STRUCTURE FOR DIETRICHSON DRILLING

Named Insured:	Dietrichson Drilling
Insurance Company:	Pacific All Risk Insurance Co.
Per-Occurrence Layer:	400,000 xs 600,000
Stop Loss Layer:	5,000,000 xs 3,000,000
Allocated expenses included in the definition of "loss"	

TABLE 2.2

SEVERITY DISTRIBUTION FOR DIETRICHSON DRILLING

Probability	Loss Amount	Excess Loss
0.00%	0	0
37.80%	200,000	0
23.50%	400,000	0
14.60%	600,000	0
9.10%	800,000	200,000
15.00%	1,000,000	400,000
Average	480,000	78,200

stop loss layer with no problem, but you recognize that there is likely to be a strong dependence between the results of the two covers and you want to reflect this in your pricing.

We will consider a simplified version of this problem. First, we will assume that the loss distribution can be reasonably approximated using only a five-point discretized severity distribution. In practice, a curve of more than a hundred points would be needed in order to accurately capture the true shape. For our example, the simpler distribution shown in Table 2.2 will be used.

Consistent with this loss distribution, our average severity is estimated to be 480,000 and the average in the 400,000 xs 600,000 layer is 78,200.

TABLE 2.3

SINGLE CLAIM PRIMARY & EXCESS LOSS BIVARIATE DISTRIBUTION

	Loss Excess of 600,000					
Loss Capped at 600,000	0	200,000	400,000	600,000	800,000	1,000,000
0	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
200,000	**37.80%**	0.00%	0.00%	0.00%	0.00%	0.00%
400,000	**23.50%**	0.00%	0.00%	0.00%	0.00%	0.00%
600,000	**14.60%**	**9.10%**	**15.00%**	0.00%	0.00%	0.00%
800,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
1,000,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
1,200,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
1,400,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
1,600,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
1,800,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
2,000,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%

We have also estimated that the expected number of claims is 5, with a variance of 6, and the frequency will be modeled using a Negative Binomial[3] distribution. The overall loss pick is therefore 2,400,000 ($5 \times 480{,}000$). Our aggregate model calculates expected losses of 123,529 in the proposed stop loss layer above 3,000,000.

The first step in calculating the overall loss distribution is to create a bivariate severity distribution of primary and excess losses. This is shown in Table 2.3.

From Table 2.3, we can observe a strong dependence structure between the primary and excess losses: we can have an excess loss only if the primary 600,000 retention is hit.

This bivariate severity matrix becomes the input for the FFT model, and may be denoted M_x. The matrix of the aggregate distribution may be denoted M_z and is produced using the two-

[3] See Appendix D for details on the Negative Binomial distribution and its Probability Generating Function.

dimensional Fast Fourier Transform calculation:

$$M_z = \text{IFFT}(\text{PGF}(\text{FFT}(M_x))), \qquad \text{and} \tag{2.1}$$

$$\text{PGF}(t) = (1.2 - .2t)^{-25}. \tag{2.2}$$

For the bivariate matrix M_x shown in Table 2.3, the resulting M_z is given in Table 2.4.

An additional step is needed in order to calculate the estimated results in the stop loss layer above 3,000,000. For that calculation, the rows of Table 2.4 for all amounts 3,000,000 or less are summed to compute the probabilities of no excess-of-aggregate losses. The remaining rows are intact but the row labels are reduced by 3,000,000. The result is Table 2.5.

From Table 2.5, several statistics of interest can be calculated.[4] The expected loss to the stop loss layer is 123,529 and the probability that the stop loss is hit is 15.08%. The average loss amount conditional upon the stop loss being hit is 819,210.

More dramatic from a risk management perspective is the dependence between the per-occurrence and stop loss covers. The expected loss to the per-occurrence layer is 391,000 ($5 \times$ 78,200), but this increases to 830,334 when we include only the scenarios in which the stop loss is also hit. This dependence needs to be considered in the decision to write the contract: *on average, when the stop loss is hit we will also be paying about twice the expected amount in the per-occurrence layer*.

The two-dimensional matrix shown in Table 2.5 can be used to verify the expected loss pricing for either coverage individually. The probabilities associated with the stop loss program are found by summing across rows; the probabilities associated with the per-occurrence excess layer are found by summing down

[4]The probabilities for the aggregate distribution extend beyond the rows and columns actually displayed.

TABLE 2.4

AGGREGATE PRIMARY & AGGREGATE EXCESS LOSS BIVARIATE DISTRIBUTION

	Loss Excess of 600,000						
Loss Capped at 600,000	0	200,000	400,000	600,000	800,000	1,000,000	1,200,000
0	1.05%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
200,000	1.65%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
400,000	2.38%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
600,000	3.09%	0.40%	0.66%	0.00%	0.00%	0.00%	0.00%
800,000	3.34%	0.65%	1.07%	0.00%	0.00%	0.00%	0.00%
1,000,000	3.39%	0.96%	1.58%	0.00%	0.00%	0.00%	0.00%
1,200,000	3.22%	1.27%	2.16%	0.26%	0.21%	0.00%	0.00%
1,400,000	2.86%	1.40%	2.44%	0.44%	0.36%	0.00%	0.00%
1,600,000	2.43%	1.45%	2.59%	0.66%	0.54%	0.00%	0.00%
1,800,000	1.97%	1.40%	2.57%	0.90%	0.78%	0.09%	0.05%
2,000,000	1.54%	1.26%	2.38%	1.02%	0.92%	0.15%	0.08%
2,200,000	1.17%	1.09%	2.12%	1.08%	1.01%	0.24%	0.13%
2,400,000	0.86%	0.90%	1.80%	1.08%	1.05%	0.33%	0.20%
2,600,000	0.62%	0.72%	1.47%	1.00%	1.01%	0.38%	0.24%
2,800,000	0.43%	0.55%	1.16%	0.88%	0.93%	0.42%	0.27%
3,000,000	0.29%	0.41%	0.89%	0.75%	0.82%	0.43%	0.29%
3,200,000	0.20%	0.30%	0.66%	0.61%	0.70%	0.41%	0.29%
3,400,000	0.13%	0.21%	0.48%	0.48%	0.57%	0.37%	0.28%
3,600,000	0.08%	0.15%	0.34%	0.37%	0.45%	0.32%	0.25%
3,800,000	0.05%	0.10%	0.24%	0.27%	0.35%	0.27%	0.22%
4,000,000	0.03%	0.07%	0.16%	0.20%	0.26%	0.22%	0.19%
4,200,000	0.02%	0.04%	0.11%	0.14%	0.19%	0.17%	0.15%
4,400,000	0.01%	0.03%	0.07%	0.10%	0.14%	0.13%	0.12%
4,600,000	0.01%	0.02%	0.05%	0.07%	0.10%	0.10%	0.09%
4,800,000	0.00%	0.01%	0.03%	0.04%	0.07%	0.07%	0.07%
5,000,000	0.00%	0.01%	0.02%	0.03%	0.04%	0.05%	0.05%
5,200,000	0.00%	0.00%	0.01%	0.02%	0.03%	0.03%	0.04%
5,400,000	0.00%	0.00%	0.01%	0.01%	0.02%	0.02%	0.03%
5,600,000	0.00%	0.00%	0.00%	0.01%	0.01%	0.02%	0.02%
5,800,000	0.00%	0.00%	0.00%	0.00%	0.01%	0.01%	0.01%
6,000,000	0.00%	0.00%	0.00%	0.00%	0.01%	0.01%	0.01%
6,200,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.01%
6,400,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
6,600,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
6,800,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
7,000,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
7,200,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
7,400,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
7,600,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
7,800,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
8,000,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%

TABLE 2.5

AGGREGATE PRIMARY EXCESS & AGGREGATE EXCESS LOSS BIVARIATE DISTRIBUTION

Stop Loss	Loss Excess of 600,000 0	200,000	400,000	600,000	800,000	1,000,000	1,200,000
0	30.28%	12.45%	22.89%	8.07%	7.65%	2.04%	1.26%
200,000	0.20%	0.30%	0.66%	0.61%	0.70%	0.41%	0.29%
400,000	0.13%	0.21%	0.48%	0.48%	0.57%	0.37%	0.28%
600,000	0.08%	0.15%	0.34%	0.37%	0.45%	0.32%	0.25%
800,000	0.05%	0.10%	0.24%	0.27%	0.35%	0.27%	0.22%
1,000,000	0.03%	0.07%	0.16%	0.20%	0.26%	0.22%	0.19%
1,200,000	0.02%	0.04%	0.11%	0.14%	0.19%	0.17%	0.15%
1,400,000	0.01%	0.03%	0.07%	0.10%	0.14%	0.13%	0.12%
1,600,000	0.01%	0.02%	0.05%	0.07%	0.10%	0.10%	0.09%
1,800,000	0.00%	0.01%	0.03%	0.04%	0.07%	0.07%	0.07%
2,000,000	0.00%	0.01%	0.02%	0.03%	0.04%	0.05%	0.05%
2,200,000	0.00%	0.00%	0.01%	0.02%	0.03%	0.03%	0.04%
2,400,000	0.00%	0.00%	0.01%	0.01%	0.02%	0.02%	0.03%
2,600,000	0.00%	0.00%	0.00%	0.01%	0.01%	0.02%	0.02%
2,800,000	0.00%	0.00%	0.00%	0.00%	0.01%	0.01%	0.01%
3,000,000	0.00%	0.00%	0.00%	0.00%	0.01%	0.01%	0.01%
3,200,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.01%
3,400,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
3,600,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
3,800,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
4,000,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
4,200,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
4,400,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
4,600,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
4,800,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
5,000,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%

columns. By summing across rows or down columns, we calculate the marginal distributions.

In order to calculate the distribution of the sum of the two coverages combined, we sum the probabilities along each lower-left to upper-right diagonal. Table 2.6 shows this calculation.

TABLE 2.6

PROBABILITIES FOR AGGREGATE PRIMARY EXCESS PLUS AGGREGATE EXCESS LOSS

Loss & ALAE	Probability	Calculation
0	30.28%	= 30.28%
200,000	12.64%	= 0.20% + 12.45%
400,000	23.31%	= 0.13% + 0.30% + 22.89%
600,000	9.02%	= 0.08% + 0.21% + 0.66% + 8.07%
800,000	8.94%	= 0.05% + 0.15% + 0.48% + 0.61% + 7.65%
⋮	⋮	⋮

3. DISTRIBUTION FOR LOSS ONLY SUBJECT TO AGGREGATE LIMIT PLUS UNLIMITED ALLOCATED LOSS ADJUSTMENT EXPENSE (ALAE)

Our insured, Dietrichson Drilling, requests a general liability policy on a traditional guaranteed cost basis. Our company, Pacific All Risk Insurance Company, is willing to offer a standard policy form with a 1,000,000 per-occurrence limit and a 2,000,000 general policy aggregate.

Both the per-occurrence limit and the general aggregate limit apply to the indemnity loss only. All defense costs and associated expenses (allocated loss adjustment expense—ALAE) are covered in addition to these limits. The Pacific All Risk policy is summarized in Table 3.1. The loss distribution is approximated in Table 3.2.

As the Pacific All Risk actuary, you have been asked to estimate the aggregate distribution of the sum of the loss and ALAE combined. The first step in calculating the overall loss distribution is to assemble the bivariate severity distribution of loss and ALAE. This is shown in Table 3.3.

For Dietrichson Drilling, we believe that there will be a strong dependence between loss and ALAE; larger losses are generally

TABLE 3.1

POLICY STRUCTURE FOR DIETRICHSON DRILLING

Named Insured:	Dietrichson Drilling
Insurance Company:	Pacific All Risk Insurance Co.
Per-Occurrence Limit:	1,000,000
General Aggregate Limit:	2,000,000
Allocated expenses paid in addition to loss	

TABLE 3.2

SEVERITY DISTRIBUTION FOR DIETRICHSON DRILLING

Probability	Loss Amount
10.00%	0
45.00%	200,000
9.00%	400,000
9.00%	600,000
9.00%	800,000
18.00%	1,000,000
Average	432,000
Average ALAE %	37.29%

assumed to have larger dollars of associated expenses. The numbers in Table 3.3 are for illustration only, but were selected to demonstrate such a dependence.

The table is constructed such that the loss severity curve does not extend beyond the 1,000,000 per-occurrence limit, whereas the ALAE curve does not have an explicit cap. By convention, we are also including closed-without-pay claims in this analysis, at least to the extent that they contribute ALAE.

This bivariate severity matrix becomes the input for the FFT model, and will again be denoted as M_x. The matrix of aggregate

TABLE 3.3

SINGLE CLAIM LOSS & ALAE BIVARIATE DISTRIBUTION

	ALAE						
Loss Amount	0	200,000	400,000	600,000	800,000	1,000,000	1,200,000
0	8.39%	1.47%	0.13%	0.01%	0.00%	0.00%	0.00%
200,000	27.98%	13.29%	3.16%	0.50%	0.06%	0.01%	0.00%
400,000	4.15%	3.21%	1.25%	0.32%	0.06%	0.01%	0.00%
600,000	3.07%	3.30%	1.77%	0.64%	0.17%	0.04%	0.01%
800,000	2.28%	3.13%	2.15%	0.99%	0.34%	0.09%	0.02%
1,000,000	3.37%	5.65%	4.73%	2.64%	1.11%	0.37%	0.10%
1,200,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
1,400,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
1,600,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
1,800,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
2,000,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%

distributions M_z is again given by the formula:

$$M_z = \mathrm{IFFT}(\mathrm{PGF}(\mathrm{FFT}(M_x))), \qquad \text{and} \tag{3.1}$$

$$\mathrm{PGF}(t) = (2 - t)^{-4}. \tag{3.2}$$

The frequency distribution is assumed to be Negative Binomial, with a mean of 4 and a variance of 8.

The final matrix of aggregate distributions is shown in Table 3.4. In order to cap the loss-only exposure at the 2,000,000 general aggregate, we sum the probabilities for losses above 2,000,000 into a single row. The result is Table 3.5. Finally, we can create a single distribution from this matrix by summing along each lower-left to upper-right diagonal to obtain Table 3.6.

It is also instructive to show a graph of the distribution of the combined loss and ALAE both before and after the general aggregate cap. In Graph 3.1 we can see that the "tail" of the cumulative distribution is greatly reduced by imposing a 2,000,000 general aggregate. However, we note that there is still a non-

TABLE 3.4

AGGREGATE LOSS & AGGREGATE ALLOCATED LOSS ADJUSTMENT EXPENSE JOINT DISTRIBUTION

Loss Amount	ALAE 0	200,000	400,000	600,000	800,000	1,000,000	1,200,000
0	7.42%	0.23%	0.02%	0.00%	0.00%	0.00%	0.00%
200,000	4.33%	2.23%	0.59%	0.11%	0.02%	0.00%	0.00%
400,000	2.22%	2.10%	1.01%	0.33%	0.08%	0.02%	0.00%
600,000	1.41%	1.82%	1.20%	0.54%	0.18%	0.05%	0.01%
800,000	1.06%	1.71%	1.40%	0.77%	0.33%	0.11%	0.03%
1,000,000	1.11%	2.09%	1.98%	1.27%	0.62%	0.25%	0.08%
1,200,000	0.69%	1.59%	1.84%	1.42%	0.83%	0.39%	0.16%
1,400,000	0.40%	1.07%	1.45%	1.32%	0.90%	0.50%	0.23%
1,600,000	0.25%	0.75%	1.15%	1.19%	0.92%	0.57%	0.30%
1,800,000	0.16%	0.56%	0.95%	1.08%	0.93%	0.64%	0.37%
2,000,000	0.12%	0.44%	0.82%	1.02%	0.96%	0.72%	0.46%
2,200,000	0.07%	0.30%	0.61%	0.84%	0.87%	0.72%	0.50%
2,400,000	0.04%	0.19%	0.43%	0.65%	0.73%	0.66%	0.50%
2,600,000	0.03%	0.13%	0.31%	0.50%	0.60%	0.59%	0.48%
2,800,000	0.02%	0.09%	0.22%	0.38%	0.50%	0.52%	0.45%
3,000,000	0.01%	0.06%	0.16%	0.30%	0.41%	0.46%	0.42%
3,200,000	0.01%	0.04%	0.11%	0.22%	0.32%	0.38%	0.37%
3,400,000	0.00%	0.02%	0.07%	0.15%	0.24%	0.30%	0.32%
3,600,000	0.00%	0.01%	0.05%	0.11%	0.18%	0.24%	0.27%
3,800,000	0.00%	0.01%	0.03%	0.08%	0.13%	0.19%	0.22%
4,000,000	0.00%	0.01%	0.02%	0.05%	0.10%	0.14%	0.18%

remote probability of loss even above 3,000,000, due to the inclusion of ALAE on an unlimited basis.

4. DYNAMIC FINANCIAL ANALYSIS

As the actuary for Pacific All Risk, you have now completed your pricing work for individual insurance contracts. As a reward for your hard work, you have been rotated to the actuarial team that runs the company's Dynamic Financial Analysis (DFA) model, called Pacific Enterprise Risk Model (PERM).

TABLE 3.5

AGGREGATE LOSS CAPPED AT 2,000,000 & AGGREGATE ALLOCATED LOSS ADJUSTMENT EXPENSE JOINT DISTRIBUTION

	ALAE						
Loss Amount	0	200,000	400,000	600,000	800,000	1,000,000	1,200,000
0	7.42%	0.23%	0.02%	0.00%	0.00%	0.00%	0.00%
200,000	4.33%	2.23%	0.59%	0.11%	0.02%	0.00%	0.00%
400,000	2.22%	2.10%	1.01%	0.33%	0.08%	0.02%	0.00%
600,000	1.41%	1.82%	1.20%	0.54%	0.18%	0.05%	0.01%
800,000	1.06%	1.71%	1.40%	0.77%	0.33%	0.11%	0.03%
1,000,000	1.11%	2.09%	1.98%	1.27%	0.62%	0.25%	0.08%
1,200,000	0.69%	1.59%	1.84%	1.42%	0.83%	0.39%	0.16%
1,400,000	0.40%	1.07%	1.45%	1.32%	0.90%	0.50%	0.23%
1,600,000	0.25%	0.75%	1.15%	1.19%	0.92%	0.57%	0.30%
1,800,000	0.16%	0.56%	0.95%	1.08%	0.93%	0.64%	0.37%
2,000,000	0.31%	1.31%	2.88%	4.41%	5.28%	5.31%	4.73%
2,200,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%

TABLE 3.6

PROBABILITIES FOR LIMITED LOSS PLUS ALAE

Combined Loss + ALAE	Probability	Calculation
0	7.42%	= 7.42%
200,000	4.56%	= 4.33% + 0.23%
400,000	4.47%	= 2.22% + 2.23% + 0.02%
600,000	4.09%	= 1.41% + 2.10% + 0.59% + 0.00%
800,000	3.99%	= 1.06% + 1.82% + 1.01% + 0.11% + 0.00%
⋮	⋮	⋮

The goal of the PERM team is to model the distribution of results for Pacific All Risk Insurance Company as a whole. Included in this analysis is sensitivity testing for interest rates and various complex reinsurance structures. The PERM is a giant simulation model that needs to be parameterized for the business actually written.

GRAPH 3.1

CUMULATIVE DISTRIBUTION FUNCTIONS FOR CAPPED AND UNCAPPED LOSS & ALAE

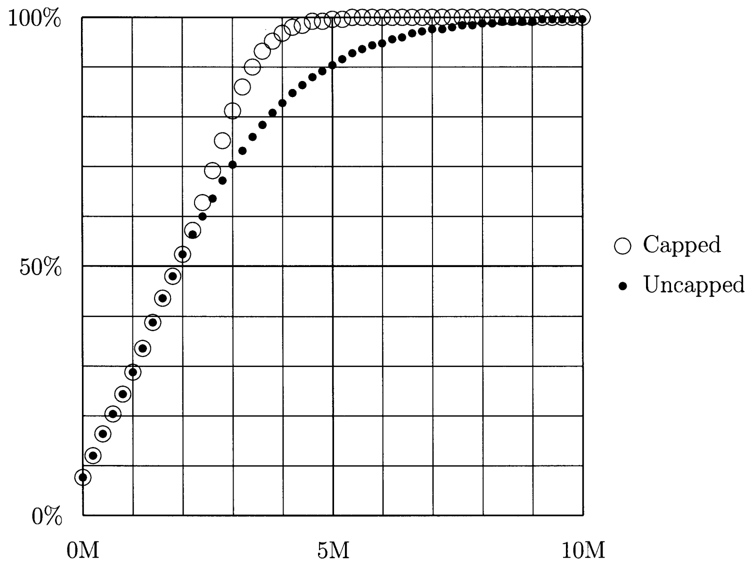

A simplification made in the PERM is that the model separately simulates an aggregate value for all “small” losses and then simulates individual “large” losses. A truncation point of 1,000,000 has been selected for segregating large from small losses.

An early version of the PERM made the assumption that the small and large losses are independent. That is, the small and large losses were simulated separately and then the results were summed. However, this independence assumption was found to be false, resulting in understated variability and unrealistically low probabilities in the tail of the combined distribution.

In fact, the aggregate distributions of the small and large losses are generally not independent. If a single frequency distribution is used to generate the overall number of losses, N, then the covariance[5] can be written explicitly

$$\mathrm{Cov}(S,L) = p\mu_S(1-p)\mu_L(\sigma_N^2 - \mu_N), \tag{4.1}$$

where

S = aggregate small losses,

L = aggregate large losses,

μ_S = conditional mean of small claim size,

μ_L = conditional mean of large claim size,

p = probability that a given claim is small,

σ_N^2 = variance of the claim counts, and

μ_N = mean of the claim counts.

The sign of the covariance term is driven by the claim count distribution. For the commonly used Negative Binomial this is positive; for the Poisson it is zero.[6] Equation (4.1) is derived in Appendix B.

In order to model the losses for Pacific All Risk, we begin by approximating the total loss distribution with a few discrete points (Table 4.1). As in the previous examples, a five-point distribution is used here, but would need to be expanded to a greater number of points in a more realistic application.

This single severity curve is then reconfigured into Table 4.2, a bivariate matrix M_x. The first column defines the severity of the "small" loss distribution. The first row is a single point containing the probability of a "large" loss.

This format is a bit different than the previous examples, since the vertical and horizontal axes are in different units: the vertical

[5] Sundt shows a more general formula in [8].

[6] In the case of the Poisson it can be shown that the large and small claims are actually independent.

TABLE 4.1

SEVERITY DISTRIBUTION

Probability	Loss Amount
0.00%	0
43.80%	200,000
24.60%	400,000
13.80%	600,000
7.80%	800,000
10.00%	1,000,000
Average	431,200

TABLE 4.2

SINGLE CLAIM SMALL LOSS & LARGE COUNTS JOINT DISTRIBUTION

	Large Loss Counts					
Small Loss	0	1	2	3	4	5
0	0.00%	**10.00%**	0.00%	0.00%	0.00%	0.00%
200,000	**43.80%**	0.00%	0.00%	0.00%	0.00%	0.00%
400,000	**24.60%**	0.00%	0.00%	0.00%	0.00%	0.00%
600,000	**13.80%**	0.00%	0.00%	0.00%	0.00%	0.00%
800,000	**7.80%**	0.00%	0.00%	0.00%	0.00%	0.00%
1,000,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
1,200,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
1,400,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
1,600,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
1,800,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
2,000,000	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%

in dollars and the horizontal in counts. This illustrates the flexibility in the FFT technique to allow for different scale factors for the two dimensions.

For a frequency distribution, we use a Negative Binomial with mean 10 and variance 20. For an actual insurance company, the overall frequency is likely to be much higher, but we continue with this simplified assumption for clarity. The aggregate distribution matrix M_z is again given by the expression:

$$M_z = \mathrm{IFFT}(\mathrm{PGF}(\mathrm{FFT}(M_x))), \qquad \text{and} \tag{4.2}$$

$$\mathrm{PGF}(t) = (2-t)^{-10}. \tag{4.3}$$

The resulting aggregate distribution matrix M_z is in Table 4.3. Like the original bivariate severity, this matrix has units in dollars for the "small" losses, and counts for the "large" losses. The marginal distribution for the aggregate small losses is found by summing the probabilities in each row.

The simulation procedure first simulates an aggregate amount for the "small" losses, and then finds a conditional frequency distribution for the "large" loss counts. The conditional large loss frequency distributions are created by rescaling each row of M_z to total 100%. This is shown in Table 4.4.

The conditional matrix shown in Table 4.4 is also instructive in itself, because it clearly shows the dependence between large and small losses. Simply put, an increase in frequency means more losses in both the large and small categories.

The final simulation procedure for the PERM is then:

- simulate the aggregate dollars of small losses out of its marginal distribution;
- simulate the number of large losses from the corresponding conditional frequency distribution;
- simulate a severity amount for each of the large losses.

This procedure allows us to efficiently simulate losses without the need to individually simulate every small loss, and at the same time preserves the dependence structure between the large and small losses.

TABLE 4.3

AGGREGATE CLAIM SMALL LOSS & LARGE COUNTS BIVARIATE DISTRIBUTION

Small Loss	Large Loss Counts 0	1	2	3	4	5	6
0	0.10%	0.05%	0.01%	0.00%	0.00%	0.00%	0.00%
200,000	0.21%	0.12%	0.04%	0.01%	0.00%	0.00%	0.00%
400,000	0.38%	0.22%	0.07%	0.02%	0.00%	0.00%	0.00%
600,000	0.58%	0.36%	0.12%	0.03%	0.01%	0.00%	0.00%
800,000	0.82%	0.53%	0.18%	0.05%	0.01%	0.00%	0.00%
1,000,000	1.07%	0.71%	0.26%	0.07%	0.01%	0.00%	0.00%
1,200,000	1.31%	0.91%	0.34%	0.09%	0.02%	0.00%	0.00%
1,400,000	1.54%	1.11%	0.43%	0.12%	0.03%	0.00%	0.00%
1,600,000	1.74%	1.30%	0.52%	0.15%	0.03%	0.01%	0.00%
1,800,000	1.90%	1.46%	0.60%	0.18%	0.04%	0.01%	0.00%
2,000,000	2.02%	1.60%	0.68%	0.20%	0.05%	0.01%	0.00%
2,200,000	2.09%	1.71%	0.75%	0.23%	0.06%	0.01%	0.00%
2,400,000	2.12%	1.78%	0.80%	0.25%	0.06%	0.01%	0.00%
2,600,000	2.11%	1.82%	0.84%	0.27%	0.07%	0.02%	0.00%
2,800,000	2.06%	1.83%	0.86%	0.29%	0.08%	0.02%	0.00%
3,000,000	1.98%	1.81%	0.87%	0.30%	0.08%	0.02%	0.00%
3,200,000	1.88%	1.75%	0.87%	0.30%	0.08%	0.02%	0.00%
3,400,000	1.76%	1.68%	0.85%	0.30%	0.09%	0.02%	0.00%
3,600,000	1.62%	1.59%	0.82%	0.30%	0.09%	0.02%	0.00%
3,800,000	1.48%	1.49%	0.79%	0.29%	0.09%	0.02%	0.00%
4,000,000	1.34%	1.38%	0.74%	0.28%	0.08%	0.02%	0.00%
4,200,000	1.20%	1.26%	0.70%	0.27%	0.08%	0.02%	0.00%
4,400,000	1.06%	1.14%	0.64%	0.25%	0.08%	0.02%	0.00%
4,600,000	0.94%	1.03%	0.59%	0.24%	0.08%	0.02%	0.00%
4,800,000	0.82%	0.91%	0.54%	0.22%	0.07%	0.02%	0.00%
5,000,000	0.71%	0.81%	0.48%	0.20%	0.07%	0.02%	0.00%

5. CONCLUSION

Aggregating a bivariate severity distribution is a useful technique. Two severity components are separately aggregated while preserving their dependence structure. This technique can be applied when pricing a policy with a per-occurrence retention and a stop loss on the aggregate retention. It can also be applied more

TABLE 4.4

CONDITIONAL DISTRIBUTIONS OF LARGE COUNTS GIVEN AGGREGATE SMALL LOSSES

Small Loss	Large Loss Counts 0	1	2	3	4	5	6
0	59.87%	29.94%	8.23%	1.65%	0.27%	0.04%	0.00%
200,000	56.88%	31.28%	9.39%	2.03%	0.36%	0.05%	0.01%
400,000	54.91%	32.07%	10.18%	2.33%	0.43%	0.07%	0.01%
600,000	53.26%	32.68%	10.87%	2.60%	0.50%	0.08%	0.01%
800,000	51.80%	33.17%	11.49%	2.85%	0.57%	0.10%	0.01%
1,000,000	50.37%	33.62%	12.11%	3.12%	0.64%	0.11%	0.02%
1,200,000	49.03%	34.01%	12.70%	3.39%	0.72%	0.13%	0.02%
1,400,000	47.77%	34.34%	13.26%	3.65%	0.80%	0.15%	0.02%
1,600,000	46.55%	34.63%	13.81%	3.92%	0.89%	0.17%	0.03%
1,800,000	45.38%	34.87%	14.34%	4.19%	0.97%	0.19%	0.03%
2,000,000	44.26%	35.09%	14.86%	4.47%	1.07%	0.22%	0.04%
2,200,000	43.17%	35.26%	15.37%	4.74%	1.16%	0.24%	0.04%
2,400,000	42.12%	35.41%	15.86%	5.02%	1.26%	0.27%	0.05%
2,600,000	41.10%	35.53%	16.34%	5.31%	1.37%	0.30%	0.06%
2,800,000	40.11%	35.62%	16.80%	5.59%	1.47%	0.33%	0.06%
3,000,000	39.16%	35.69%	17.25%	5.88%	1.58%	0.36%	0.07%
3,200,000	38.22%	35.73%	17.69%	6.17%	1.70%	0.39%	0.08%
3,400,000	37.32%	35.75%	18.12%	6.46%	1.82%	0.43%	0.09%
3,600,000	36.44%	35.75%	18.54%	6.75%	1.94%	0.47%	0.10%
3,800,000	35.58%	35.73%	18.94%	7.05%	2.07%	0.51%	0.11%
4,000,000	34.74%	35.69%	19.33%	7.34%	2.19%	0.55%	0.12%
4,200,000	33.93%	35.63%	19.71%	7.64%	2.33%	0.59%	0.13%
4,400,000	33.14%	35.56%	20.08%	7.94%	2.46%	0.64%	0.14%
4,600,000	32.37%	35.47%	20.44%	8.24%	2.61%	0.69%	0.16%
4,800,000	31.61%	35.36%	20.78%	8.54%	2.75%	0.74%	0.17%
5,000,000	30.88%	35.25%	21.12%	8.83%	2.90%	0.79%	0.19%

generally. The two random variables can be different items such as dollars and counts.

In this paper we aggregate the bivariate distribution using the FFT, but it is possible to do this with the continuous Fourier Transform or simulation. Sundt [8] shows that this can be done with recursive techniques. It may sometimes be preferable to utilize a mix of techniques.

This technique can be extended to n dimensions by developing a multivariate distribution M_x. With the claim count PGF and an n-dimensional FFT, the aggregate multivariate array M_z is obtained as

$$M_z = \text{IFFT}(\text{PGF}(\text{FFT}(M_x))). \tag{5.1}$$

REFERENCES

[1] Heckman, Philip E., and Glenn G. Meyers, "The Calculation of Aggregate Loss Distributions from Claim Severity and Claim Count Distributions," *PCAS* LXX, 1983, pp. 22–61.

[2] Ihaka, Ross, and Robert Gentleman, "R: A Language for Data Analysis and Graphics," *Journal of Computational and Graphical Statistics* 5, 3, 1996, pp. 299–314.

[3] Klugman, Stuart A., Harry H. Panjer, and E. Willmot Gordon, *Loss Models: From Data to Decisions*, New York, John Wiley & Sons, Inc., 1998.

[4] Panjer, Harry H., "Recursive Evaluation of a Family of Compound Distributions," *ASTIN Bulletin*, 12, 1, 1981, pp. 22–26.

[5] Papush, Dmitry E., Gary S. Patrik, and Felix Podgaits, "Approximations of the Aggregate Loss Distribution," *CAS Forum*, Winter 2001, pp. 175–186.

[6] Press, William H., Brian P. Flannery, Saul A. Teukolsky, and William T. Vetterling, *Numerical Recipes in C: The Art of Scientific Computing*, Cambridge University Press, 1988.

[7] Robertson, John P., "The Computation of Aggregate Loss Distributions," *PCAS* LXXIX, 1992, pp. 57–133.

[8] Sundt, Bjørn, "On Multivariate Panjer Recursions," *ASTIN Bulletin*, 29, 1, 1999, pp. 29–45.

[9] Venter, Gary G., "Transformed Beta and Gamma Distributions and Aggregate Losses," *PCAS* LXX, 1983, pp. 156–193.

[10] Wang, Shaun S., "Aggregation of Correlated Risk Portfolios: Models and Algorithms," *PCAS* LXXXV, 1998, pp. 848–939.

APPENDIX A

SAMPLE TWO-DIMENSIONAL FAST FOURIER TRANSFORM USING R

It is convenient to compute FFTs using preprogrammed software. An excellent piece of software that includes FFT functions is based on the S language and is publicly available for free. It is called "R" [2]. Versions of R for various operating systems can be found by following 'http://cran.r-project.org/'. R is copyrighted software made publicly available under the GNU General Public License which is available at 'http://www.gnu.org/copyleft/gpl.html'. The FFT function is also available in commercial software packages, e.g., MATLAB and S-Plus.

A listing from a session with R shows how easy it is to compute two-dimensional FFTs. Lines typed by the user begin with ">". The inverse of a matrix *M* is obtained with "`fft(M,T)/n`," where *n* is the number of elements in the matrix.

```
>ms<-matrix(c(.4,0,0,.3,.3,0,0,0,0),3,3,byrow=T)
>ms
[,1] [,2] [,3]
[1,] 0.4 0.0 0
[2,] 0.3 0.3 0
[3,] 0.0 0.0 0
>f<-fft(ms)
>f
[,1] [,2] [,3]
[1,] 1.0+0.0000000i 0.55-0.2598076i 0.55+0.2598076i
[2,] 0.1-0.5196152i 0.10+0.0000000i 0.55-0.2598076i
[3,] 0.1+0.5196152i 0.55+0.2598076i 0.10+0.0000000i
>f*f
[,1] [,2] [,3]
[1,] 1.00+0.0000000i 0.235-0.2857884i 0.235+0.2857884i
[2,] -0.26-0.1039230i 0.010+0.0000000i 0.235-0.2857884i
[3,] -0.26+0.1039230i 0.235+0.2857884i 0.010+0.0000000i
>ma<-fft(f*f,T)/9
>ma
```

```
[,1] [,2] [,3]
[1,] 0.16+0i 1.652685e-18+0i 2.301894e-17+0i
[2,] 0.24+0i 2.400000e-01+0i 2.467162e-17+0i
[3,] 0.09+0i 1.800000e-01+0i 9.000000e-02+0i
```

For those wishing to program their own algorithms, see [6]. Note that, when the object to be transformed consists only of real numbers, there are symmetries that can be used to decrease the amount of computing required. Also note that many software packages, including R, define the FFT as $\mathrm{FFT}(x)_k = \sum_{j=0}^{n-1} \exp(-2\pi ijk/n)$, using a negative exponent instead of a positive one as we have in Equation (C.1). The corresponding inverse is $\mathrm{IFFT}(\tilde{x})_k = (1/n)\sum_{j=0}^{n-1} \exp(2\pi ijk/n)$. The reader wishing to verify his code with a package like R should use the "negative" sign convention.

APPENDIX B

CORRELATION OF LARGE AND SMALL LOSSES

Consider the Collective Risk Model with aggregate losses represented by the sum of individual claims

$$Z = X_1 + \cdots + X_N. \tag{B.1}$$

The X_i are independent and identically distributed (iid) random variables denoting claim sizes. Claim counts are denoted by the random variable N, which is independent from each X_i. It is further assumed that the first moment of X_i is finite and that the second moment of N is finite.

Let T denote the threshold for distinguishing between small claims and large claims; i.e., X_i is small if $X_i \leq T$. Define a small loss indicator, $I_i = 1$ for $X_i \leq T$ and 0 otherwise. Then we have small aggregate losses

$$Z_S = X_1 I_1 + \cdots + X_N I_N, \tag{B.2}$$

and large aggregate losses

$$Z_L = X_1(1 - I_1) + \cdots + X_N(1 - I_N). \tag{B.3}$$

Let p be the probability that $X_i \leq T$. Denote the conditional means for small and large claim sizes with

$$\mu_S = \mathrm{E}[X_i | X_i \leq T], \qquad \text{and} \tag{B.4}$$

$$\mu_L = \mathrm{E}[X_i | X_i > T]. \tag{B.5}$$

Denote the claim count mean and variance with

$$\mu_N = \mathrm{E}[N], \qquad \text{and} \tag{B.6}$$

$$\sigma_N^2 = \mathrm{Var}[N]. \tag{B.7}$$

PROPOSITION

$$\mathrm{Cov}[Z_S, Z_L] = p\mu_S(1 - p)\mu_L(\sigma_N^2 - \mu_N). \tag{B.8}$$

Proof

$$\begin{aligned}
\mathrm{E}[Z_S Z_L] &= \mathrm{E}\left[\left(\sum_{i=1}^{N} X_i I_i\right)\left(\sum_{i=1}^{N} X_i(1-I_i)\right)\right] \\
&= \mathrm{E}_N \mathrm{E}_X \left[\left(\sum_{i=1}^{N} X_i I_i\right)\left(\sum_{i=1}^{N} X_i(1-I_i)\right)\right] \\
&= \mathrm{E}_N \mathrm{E}_X \left[\left(\sum_{i=j} X_i I_i X_j(1-I_j)\right) + \left(\sum_{j\neq i} X_i I_i X_j(1-I_j)\right)\right] \\
&= \mathrm{E}_N[N(N-1)\mathrm{E}_X[XI]\mathrm{E}_X[X(1-I)]], \\
&\qquad \text{since} \quad I_i(1-I_j) = 0 \quad \text{for} \quad i = j \\
&= \mathrm{E}_N[N(N-1)\mu_S p \mu_L(1-p)] \\
&= (\mathrm{E}(N^2) - \mu_N)\mu_S p \mu_L(1-p). \qquad \text{(B.9)}
\end{aligned}$$

$$\mathrm{E}[Z_L]\mathrm{E}[Z_S] = (\mu_N \mu_S p)(\mu_N \mu_L(1-p)) = \mu_N^2 \mu_S p \mu_L(1-p). \qquad \text{(B.10)}$$

These yield Equation (B.8), since

$$\mathrm{Cov}[Z_S, Z_L] = \mathrm{E}[Z_S Z_L] - \mathrm{E}[Z_S]\mathrm{E}[Z_L]. \qquad \text{(B.11)}$$

APPENDIX C

THE DISCRETE FOURIER TRANSFORM AND THE PROBABILITY GENERATING FUNCTION

In many insurance applications, we need to calculate an aggregate distribution in which the claim size X and the number of claims N are independent random variables. The aggregate losses are $Z = X_1 + \cdots + X_N$, where the X_k are independent and identically distributed as X. This is known as the *Collective Risk Model*.

The Fast Fourier Transform[7] (FFT) together with the Probability Generating Function (PGF) of the claim count distribution provide a convenient technique for computing the distribution of the aggregate losses.

This Appendix lists the key definitions and theorems underlying this technique. The authors recommend that the reader interested in a more comprehensive review refer to Robertson [7] and Wang [10].

C.1. Definition of FFT

We assume the claim size random variable X is discrete and describe it with an n element probability vector $x = (x_0, \ldots, x_{n-1})$, where $\text{Prob}(X = ck) = x_k$, and c is a scalar constant. For the claim count N we know the probability of each possible number of claims $\text{Prob}(N = j)$ $(j = 0, 1, \ldots)$. Let $\text{FFT}(x)$ denote the FFT of x. $\text{FFT}(x)$ is a vector with elements,

$$\tilde{x}_k = \text{FFT}(x)_k = \sum_{j=0}^{n-1} x_j \exp(2\pi i jk/n), \qquad \text{(C.1)}$$

[7]The Fast Fourier Transform is a specific implementation of the Discrete Fourier Transform. Following Wang [10] we use the term FFT for both.

where $i = \sqrt{-1}$. The FFT is also invertible; let IFFT denote its inverse.

$$x_k = \mathrm{IFFT}(\tilde{x})_k = \frac{1}{n}\sum_{j=0}^{n-1}\tilde{x}_j \exp(-2\pi i jk/n). \tag{C.2}$$

C.2. Convolution

If Z is the sum of random variables X and Y, i.e., $Z = X + Y$, then its probability vector z is known as the *convolution* of x and y and is denoted $x * y$, where

$$(x * y)_k = \sum_{l=0}^{k} x_l y_{k-l}. \tag{C.3}$$

Similarly, if Z is the sum of j independent random variables identically distributed as X, then its probability vector z is known as the jth fold convolution of x and is denoted $x^{*j} = (x_0^{*j}, \ldots, x_{n-1}^{*j})$.

It is convenient to define

$$x^{*0} = (1,0,\ldots,0). \tag{C.4}$$

The jth fold convolution can then be computed recursively for $j \geq 1$.

$$x_k^{*j} = \sum_{l=0}^{k} x_l x_{k-l}^{*j-1}. \tag{C.5}$$

C.3. Convolution Theorem for the Discrete Fourier Transform

THEOREM *Let x and y denote the probability vectors of random variables X and Y with n elements. If $x_k = y_k = 0$ for all $k \geq n/2$ when n is even and for all $k \geq (n+1)/2$ when n is odd, then*

$$\mathrm{FFT}(x * y)_k = \mathrm{FFT}(x)_k \cdot \mathrm{FFT}(y)_k. \tag{C.6}$$

For convenience we write $\mathrm{FFT}(x * y) = \mathrm{FFT}(x)\mathrm{FFT}(y)$, with the understanding that the multiplication is applied elementwise.

Extending this convention to powers, we also write $(\mathrm{FFT}(x))^j$ for $((\mathrm{FFT}(x)_0)^j, \ldots, (\mathrm{FFT}(x)_{n-1})^j)$.

Applying the inverse transform to both sides of Equation (C.6) we obtain a method for computing the convolution of two variables

$$x * y = \mathrm{IFFT}(\mathrm{FFT}(x)\mathrm{FFT}(y)). \tag{C.7}$$

Similarly, for the jth fold convolution of x

$$\mathrm{FFT}(x^{*j}) = (\mathrm{FFT}(x))^j, \tag{C.8}$$

and

$$x^{*j} = \mathrm{IFFT}(\mathrm{FFT}(x)^j). \tag{C.9}$$

C.4. Wrapping

Consider the jth fold convolution of x and note that $x_0^{*j} + \cdots + x_{n-1}^{*j}$ is not necessarily equal to 1, because n is finite. For example, suppose we have $n = 3$ and $x = (0,1,0)$. Then $x^{*3} = (0,0,0)$, since

$$\begin{aligned} x^{*1} &= (0,1,0) \\ x^{*2} &= (0 \times 0, 1 \times 0 + 0 \times 1, 0 \times 0 + 1 \times 1 + 0 \times 0) = (0,0,1) \\ x^{*3} &= (0 \times 0, 1 \times 0 + 0 \times 0, 0 \times 0 + 1 \times 0 + 0 \times 1) = (0,0,0). \end{aligned}$$

With $n = 4$ we would have $x = (0,1,0,0)$ and $x^{*3} = (0,0,0,1)$, since

$$\begin{aligned} x^{*1} &= (0,1,0,0) \\ x^{*2} &= (0 \times 0, 1 \times 0 + 0 \times 1, 0 \times 0 + 1 \times 1 + 0 \times 0, \\ &\qquad 0 \times 0 + 0 \times 1 + 1 \times 0 + 0 \times 0) \\ &= (0,0,1,0) \\ x^{*3} &= (0 \times 0, 1 \times 0 + 0 \times 0, 0 \times 0 + 1 \times 0 + 0 \times 1, \\ &\qquad 0 \times 0 + 0 \times 0 + 1 \times 1 + 0 \times 0) \\ &= (0,0,0,1). \end{aligned}$$

Robertson [7] describes the convolution in (C.3) as an *un-wrapped* convolution. Equation (C.9) returns an un-wrapped convolution if x is properly padded with zeros.

When x is not properly padded (C.9) returns a *wrapped*[8] convolution. Let $(x \tilde{*} y)$ denote a wrapped convolution; then

$$(x \tilde{*} y)_k = \sum_{j=0}^{n-1} x_j y_{(k-j) \bmod n}, \tag{C.10}$$

and

$$x_k^{\tilde{*} j} = \sum_{l=0}^{n-1} x_l x_{(k-l) \bmod n}^{\tilde{*} j-1}. \tag{C.11}$$

The wrapped convolution $x^{\tilde{*}3}$ for $n = 3$ is computed as

$$\begin{aligned}
x^{\tilde{*}1} &= (0,1,0) \\
x^{\tilde{*}2} &= ((0,1,0)\cdot(0,0,1),(0,1,0)\cdot(1,0,0),(0,1,0)\cdot(0,1,0)) \\
&= (0 \times 0 + 1 \times 0 + 0 \times 1, 0 \times 1 + 1 \times 0 + 0 \times 0, \\
&\quad 0 \times 0 + 1 \times 1 + 0 \times 0) \\
&= (0,0,1) \\
x^{\tilde{*}3} &= ((0,1,0)\cdot(0,1,0),(0,1,0)\cdot(0,0,1),(0,1,0)\cdot(1,0,0)) \\
&= (1,0,0).
\end{aligned}$$

The probability mass that is truncated with the un-wrapped convolution *wraps* with the wrapped convolution. Equation (C.9) *always* produces a wrapped convolution, but wrapped and un-wrapped convolutions are equal when x is properly padded with zeros.

[8] Robertson [7] calls this a *regular* convolution.

C.5. Definition of the Probability Generating Function

For a random variable N the Probability Generating Function (PGF) is

$$\mathrm{PGF}_N(t) = \sum_{j=0}^{\infty} t^j \mathrm{Prob}(N = j) = \mathrm{E}(t^N). \qquad \text{(C.12)}$$

The PGF for the Negative Binomial distribution is given in Appendix D.

C.6. Collective Risk Theorem for the Discrete Fourier Transform

We now show how to compute the aggregate probability vector z for the collective risk model using the FFT and the PGF of the claims count. This technique has an error term R due to wrapping which can be made arbitrarily small with sufficient zero padding of the claim size vector.

THEOREM *Suppose we have a collective risk model with claim size probability vector* $x = (x_0,\ldots,x_{n-1})$. *Let* PGF_N *be the Probability Generating Function for the claim counts* N. *Let* M *be the largest integer such that* $x_0^{*M} + \cdots + x_{n-1}^{*M} = 1$. *That is,* M *is the largest number of times one can convolute* x *and still have room for all the probability mass. Then*

$$z = \mathrm{IFFT}(\mathrm{PGF}_N(\mathrm{FFT}(x))) + R, \qquad \textit{where} \qquad \text{(C.13)}$$

$$|R_k| \le \sum_{j=M+1}^{\infty} \mathrm{Prob}(N = j). \qquad \text{(C.14)}$$

Proof

$$z = \sum_{j=0}^{\infty} x^{*j} \mathrm{Prob}(N = j). \qquad \text{(C.15)}$$

Define

$$d(j) = \mathrm{IFFT}((\mathrm{FFT}(x))^j). \qquad \text{(C.16)}$$

Then

$$z = \sum_{j=0}^{\infty} d(j)\text{Prob}(N = j) + \sum_{j=0}^{\infty} (x^{*j} - d(j))\text{Prob}(N = j). \tag{C.17}$$

Let R denote the second sum. Then

$$\begin{aligned} z &= \sum_{j=0}^{\infty} d(j)\text{Prob}(N = j) + R \\ &= \sum_{j=0}^{\infty} \text{IFFT}((\text{FFT}(x))^j)\text{Prob}(N = j) + R \\ &= \sum_{j=0}^{\infty} \text{IFFT}((\text{FFT}(x))^j \text{Prob}(N = j)) + R. \end{aligned} \tag{C.18}$$

Because IFFT is linear and continuous, we can bring it outside the summation. So,

$$\begin{aligned} z &= \text{IFFT}\left(\sum_{j=0}^{\infty} (\text{FFT}(x))^j \text{Prob}(N = j)\right) + R \\ &= \text{IFFT}(\text{PGF}_N(\text{FFT}(x))) + R. \end{aligned} \tag{C.19}$$

Now,

$$R = \sum_{j=M+1}^{\infty} (x^{*j} - d(j))\text{Prob}(N = j), \tag{C.20}$$

since $d(j) = x^{*j}$ for $j \le M$. Also, $|x_k^{*j} - d(j)_k| < 1$, since each $d(j)_k, x_k^{*j} \in [0,1]$.

Thus,

$$|R_k| \le \sum_{j=M+1}^{\infty} \text{Prob}(N = j). \tag{C.21}$$

APPENDIX D

THE NEGATIVE BINOMIAL DISTRIBUTION

The Negative Binomial distribution with parameters p and k has

$$\text{Prob}(N = j) = \frac{\Gamma(k+j)}{\Gamma(k)j!} p^k (1-p)^j, \qquad \text{and} \tag{D.1}$$

$$\text{PGF}(t) = p^k (1-(1-p)t)^{-k}. \tag{D.2}$$

The Negative Binomial mean and variance are

$$\text{Mean} = M = \frac{k(1-p)}{p}, \qquad \text{and} \tag{D.3}$$

$$\text{Variance} = V = \frac{k(1-p)}{p^2}. \tag{D.4}$$

In terms of the mean and variance the PGF is

$$\text{PGF}(t) = (V/M - (V/M - 1)t)^{M^2/(V-M)}. \tag{D.5}$$

For example, a Negative Binomial with a mean of 5 and a variance of 6 has the Probability Generating Function

$$\text{PGF}(t) = (6/5 - (6/5 - 1)t)^{5^2/(6-5)} = (1.2 - .2t)^{-25}. \tag{D.6}$$