

Risky Loss Distributions And Modeling the Loss Reserve Pay-out Tail

J. David Cummins*
University of Pennsylvania
3303 Steinberg Hall-Dietrich Hall
3620 Locust Walk
Philadelphia, PA 19104-6302
cummins@wharton.upenn.edu

James B. McDonald
Department of Economics
Brigham Young University
Provo, Utah 84602
James_McDonald@BYU.edu

Craig Merrill
Brigham Young University
Brigham Young University
678 TNRB
Provo, UT 84604
Merrill@byu.edu

October 12, 2004

Abstract: Although an extensive literature has developed on modeling the loss reserve runoff triangle, the estimation of severity distributions applicable to claims settled in specific cells of the runoff triangle has received little attention in the literature. This paper proposes the use of a very flexible probability density function, the generalized beta of the 2nd kind (GB2) to model severity distributions in the cells of the runoff triangle and illustrates the use of the GB2 based on a sample of nearly 500,000 products liability paid claims. The results show that the GB2 provides a significantly better fit to the severity data than conventional distributions such as the Weibull, Burr 12, and generalized gamma and that modeling severity by cell is important to avoid errors in estimating the riskiness of liability claims payments, especially at the longer lags.

Keywords: Loss distributions, loss reserves, generalized beta distribution, liability insurance.

JEL Classifications: C16, G22

Subject and Insurance Branch Codes: IM11, IM42, IB50

***Corresponding author.** Phone: 215-898-5644. Fax: 215-898-0310.

1. Introduction

Modeling the payout tail in long-tail lines of insurance such as general liability is an important problem in actuarial and financial modeling for the insurance industry that has received a significant amount of attention in the literature (e.g., Reid 1978; Wright 1990; Taylor 1985, 2000; Wiser, Cockley, and Gardner 2001). Modeling the payout tail is critically important in pricing, reserving, reinsurance decision making, solvency testing, dynamic financial analysis, and a host of other applications. A wide range of techniques has been developed to improve modeling accuracy and reliability. Most of the existing models focus on estimating the total claims payout in the cells of the loss runoff triangle, i.e., the variable analyzed is c_{ij} , where c_{ij} is defined as the amount of claims payment in runoff period j for accident year i . The value of c_{ij} is in turn determined by the frequency and severity of the losses in the ij -th cell of the triangle.

Although sophisticated models have been developed for estimating claim counts (frequency) and total expected payments by cell of the runoff triangle, less attention has been devoted to estimating loss severity distributions by cell. While theoretical models have been developed based on the *assumption* that claim severities by cell are gamma distributed (e.g., Mack 1991, Taylor 2000, p. 223), few empirical analyses have been conducted to determine the loss severity distributions that might be applicable to claims falling in specific cells of the runoff triangle. The objective of the present paper is to remedy this deficiency in the existing literature by conducting an extensive empirical analysis of U.S. products liability insurance paid claims. We propose the use of a flexible four-parameter distribution – the generalized beta distribution of the 2nd kind (GB2) – to model claim severities by runoff cell. This distribution is sufficiently flexible to model both heavy-tailed and light-tailed severity statistics and provides a convenient functional form for computing prices and reserve estimates. It is important to estimate loss

distributions applicable to individual cells of the runoff triangle rather than to use a single distribution applicable to all observed claims or to discount claims to present value and then fit a distribution. If the characteristics of claims settled differ significantly by settlement lag, the use of a single severity distribution can lead to severe inaccuracies in estimating expected costs, risk, and other moments of the severity distribution. This problem is likely to be especially severe for liability insurance, where claims settled at longer lags tend to be larger and more volatile.

When distributions are fit to separate years in the payout tail, the aggregate loss distribution is a mixture distribution over the yearly distributions. To explore the economic implications associated with the alternative estimates of loss distributions, we compare a single aggregate fitted distribution based on all claims for a given accident year vs. the mixture distribution, using Monte Carlo simulations. In illustrating the differences between the two models, we innovate by comparing the distributions of the discounted or economic value of claim severities rather than using undiscounted values that do not reflect the timing of payment of individual claims, thus creating *discounted severity distributions*. Thus, we provide a model (the mixture model) that not only reflects the modeling of claim severities by runoff cell but also could be used in a system designed to obtain market values of liabilities for use in fair value accounting estimation and other financial applications.

The problem of estimating claim severity distributions by cell of the runoff triangle has been previously considered by the Insurance Services Office (ISO) (1994, 1998, 2002). In ISO (1994) and (1998), a mixture of Pareto distributions was used to model loss severities by settlement lag period for products and completed operations liability losses. The two-parameter version of the Pareto was used, and the mixture consisted of two Pareto distributions. In ISO (2002), the mixed Pareto was replaced by a mixed exponential distribution, where the number of

distributions in the mixture ranged from five to eight. The ISO models do not utilize discounting or any other technique to recognize the time value of money.

Although the ISO mixture approach clearly has the potential to provide a good fit to loss severity data, we believe that there are several advantages to using a single general distribution such as the GB2 rather than a discrete mixture to model loss severity distributions by payout lag cell. The GB2 is an extremely flexible distribution that has been shown to have excellent modeling capabilities in a wide range of applications, including models of security returns and insurance claims (e.g., Bookstaber and McDonald 1987, Cummins et al. 1990, Cummins, Lewis, and Phillips 1999). It is also more natural and convenient to conduct analytical work such as price estimation and the analysis of losses by layer utilizing a single distribution rather than a mixture. The GB2 also lends itself more readily to Monte-Carlo simulation than a mixture. And, finally, the GB2 and various members of the GB2 family can be obtained analytically as general mixtures of simpler underlying distributions, so that the GB2 is in this respect already more general than a discrete mixture of Paretos or exponentials.

In addition to proposing an alternative to the ISO method for estimating severity distributions by payout lag and introducing the idea of discounted severity distributions, this paper also contributes to the existing literature by providing the first major application of the GB2 distribution to the modeling of liability insurance losses. We demonstrate that fitting a separate distribution to each year of the payout tail can lead to large differences in estimating both expected losses and the variability of losses. These differences in estimation can have a significant impact on pricing, reserving, and risk management decisions, including asset/liability management and the calculation of value at risk (VaR). We also show that the four-parameter GB2 distribution is significantly more accurate in modeling risky claim distributions than

traditional two or three-parameter distributions such as the lognormal, gamma, Weibull, or generalized gamma.

This paper builds on previous contributions in a number of excellent papers that have developed models of insurance claim severity distributions. Hogg and Klugman (1984) and Klugman, Panjer, and Willmot (1998) discuss a wide range of alternative models for loss distributions. Paulson and Faris (1985) applied the stable family of distributions, and Aiuppa (1988) considered the Pearson family as models for insurance losses. Ramlau-Hansen (1988) modeled fire, windstorm, and glass claims using the log-gamma and lognormal distributions. Cummins, et al. (1990) considered the four-parameter generalized beta of the second kind (GB2) distribution as a model for insured fire losses; and Cummins, Lewis, and Phillips (1999) use the lognormal, Burr 12, and GB2 distributions to model the severity of insured hurricane and earthquake losses. All of these papers show that the choice of distribution matters and that conventional distributions such as the lognormal and two-parameter gamma often underestimate the risk inherent in insurance claim distributions.

The paper is organized as follows: In section 2, we introduce the *GB2* family and discuss our estimation methodology. Section 3 describes the database and presents the estimated loss severity distribution results. The implications of these results are summarized in the concluding comments of section 4.

2. Statistical Models

This section reviews a family of flexible parametric probability density functions (pdf) that can be used to model insurance losses. We begin by defining a four-parameter generalized beta (GB2) probability density function, which includes many of the models considered in the prior literature as special cases. We then describe the GB2 distribution, its moments,

interpretation of parameters, and issues of estimation. This paper applies several special cases of the GB2 to explore the distribution of an extensive database on product liability claims. Previously, the GB2 has been successfully used in insurance to model fire and catastrophic property losses (Cummins, et al. 1990, Cummins, Lewis, and Phillips 1999) and has been used by a few other researchers such as Venter (1984).

2.1 The Generalized Beta of the Second Kind Distribution

The GB2 probability density function (pdf) is defined by

$$GB2(y; a, b, p, q) = \frac{|a|y^{ap-1}}{b^{ap}B(p, q)(1 + (y/b)^a)^{(p+q)}} \quad (1)$$

for $y > 0$ and zero otherwise, with b , p , and q positive, where $B(p, q)$ denotes the beta function defined by

$$B(p, q) = \int_0^1 t^{p-1} (1-t)^{q-1} dt = \frac{\Gamma(p)\Gamma(q)}{\Gamma(p+q)} \quad (2)$$

and $\Gamma(\cdot)$ denotes the gamma function. The moments of the GB2 are given by

$$E_{GB2}(y^h) = \frac{b^h B(p+h/a, q-h/a)}{B(p, q)} \quad (3)$$

The GB2 distribution includes numerous other distributions as special or limiting cases. Each special case is obtained by constraining the parameters of the more general distributions. For example, another important special case, rather a limiting case, of the generalized beta is the generalized gamma (GG)

$$\begin{aligned} GG(y; a, \beta, p) &= \lim_{q \rightarrow \infty} GB2(y; a, b = \beta q^{1/a}, p, q) \\ &= \frac{|a|y^{ap-1} e^{-(y/b)^a}}{\beta^{ap} \Gamma(p)} \end{aligned} \quad (4)$$

for $y > 0$ and zero otherwise. The moments of the generalized gamma can be expressed as

$$E_{GG}(y^h) = \beta^h \frac{\Gamma(p + h/a)}{\Gamma(p)} . \quad (5)$$

In this special case of the GB2 distribution the parameter b has been constrained to increase with q in such a way that the GB2 approaches the GG.

2.2 Interpretation of Parameters

The parameters a , b , p , and q generally determine the shape and location of the density in a complex manner. The h^{th} order moments are defined for the GG if $0 < p + h/a$ and for the GB2 if $-p < h/a < q$. Thus we see that these models permit the analysis of situations characterized by infinite variance and higher order moments. The parameter b is merely a scale parameter and depends upon the units of measurement.

Generally speaking, the larger the value of a or q , the “thinner” the tails of the density function. In fact, for “large” values of the parameter a , the probability mass of the corresponding density function becomes concentrated near the value of the parameter b . This can be verified by noting that as the parameter a increases in value the mean and variance approach b and zero, respectively. The definition of the generalized distributions permits negative values of the parameter a . This admits “inverse” distributions and in the case of the generalized gamma is called the inverse generalized gamma. Special cases of the inverse generalized gamma are used as mixing distributions in models for unobserved heterogeneity. Butler and McDonald (1987) used the GB2 as a mixture distribution.

The parameters p and q are important in determining shape. For example, for the GB2, the relative values of the parameters p and q determine the value of skewness and permit positive or negative skewness. This is in contrast to such distributions as the lognormal that are always positively skewed.

2.3 Relationships With Other Distributions

Special cases of the GB2 include the beta of the first and second kind (B1 and B2), Burr types 3 and 12 (BR3 and BR12), lognormal (LN), Weibull (W), gamma (GA), Lomax, uniform, Rayleigh, chi-square, and exponential distributions. These properties and interrelationships have been developed in other papers (e.g., McDonald, 1984, McDonald and Xu, 1995, Venter, 1984, and Cummins et al., 1990) and will not be replicated in this paper. However, since prior insurance applications have found the Burr distributions to provide excellent descriptive ability, we will formally define those pdfs:

$$\begin{aligned} BR3(y; a, b, p) &= GB2(y; a, b, p, q = 1) \\ &= \frac{|a|py^{ap-1}}{b^{ap}(1+(y/b)^a)^{p+1}} \end{aligned} \quad (6)$$

and

$$\begin{aligned} BR12(y; a, b, q) &= GB2(y; a, b, p = 1, q) \\ &= \frac{|a|qy^{a-1}}{b^a(1+(y/b)^a)^{q+1}}. \end{aligned} \quad (7)$$

Again, notice that the first line of both (6) and (7) show the relationship between the BR3 or BR12 and the GB2 distribution. The GB2 distribution includes many distributions contained in the Pearson family (see Elderton and Johnson, 1969 and Johnson and Kotz, 1970), as well as distributions such as the BR3 and BR12 which are not members of the Pearson family. Neither the Pearson nor generalized beta family nests the other.

The selection of a statistical model should be based on flexibility and ease of estimation. In numerous applications of the GB2 and its special cases, the GB2 is the best fitting four-parameter model and the BR3 and BR12 the best fitting three-parameter models.

2.4 Parameter Estimation

Methods of maximum likelihood can be applied to estimate the unknown parameters in the models discussed in the previous sections. This involves maximizing

$$l(\theta) = \sum_{t=1}^N \ln(f(y_t; \theta)) \quad (8)$$

over θ where $f(y_t; \theta)$ denotes the pdf of independent and uncensored observations of the random variable Y and N is the number of observations. In the case of censored observations the log-likelihood function becomes

$$l(\theta) = \sum_{t=1}^N [I_t \ln(f(y_t; \theta)) + (1 - I_t) \ln(1 - F(y_t; \theta))] \quad (9)$$

where $F(y_t; \theta)$ denotes the distribution function and I_t is an indicator function equal to 1 for uncensored observations and zero otherwise.¹ When I_t equals zero, i.e., a censored observation, $F(y_t; \theta)$ is evaluated at y_t equal to the policy limit plus loss adjustment expenses.

3. Estimation of Liability Severity Distributions

In this section, the methodologies described in section 2 are applied to the Insurance Services Office (ISO) closed claim paid severity data for products liability insurance. We not only fit distributions to aggregate loss data for each **accident** year but separate distributions are also fit to the claims in each cell of the payout triangle, by accident year and by settlement lag, for the years 1973 to 1986. Several distributions are used in this analysis. This section begins with a description of the database and a summary of the estimation of the severity distributions by cell of the payout triangle. The increase in risk over time and across lags is considered using

¹ All optimizations considered in this paper were performed using the programs GQOPT, obtained from Richard Quandt at Princeton, and Matlab.

means, variances, and medians. We then turn to a discussion of the estimation of the overall discounted severity distributions for each accident year using a single distribution for each year and a mixture of distributions based on the distributions applicable to the cells of the triangle.

3.1 The Database

Liability policies typically include a coverage period or accident period (usually one year) during which specified events (occurrences) are covered by the insurer. After the end of the coverage period, no new events become eligible for payment. However, the payment date for a covered event is not limited to the coverage period and may occur at any time after the date of the event.² Because of the operation of the legal liability system, payments for covered events from any given accident period extend over a long period of time after the end of the coverage period. The payout period following the coverage period is often called the runoff period or payout tail.

The database consists of products liability losses covering accident years 1973 through 1986 obtained from the Insurance Services Office (ISO). Data are on an occurrence basis, i.e., the observations represent paid and/or reserved amounts aggregated by occurrence, where an occurrence is defined as an event that gives rise to a payment or reserve. Because claim amounts are aggregated within occurrences, a single occurrence loss amount may represent payments to multiple plaintiffs for a given loss event. Claim amounts represent the total of bodily injury and property damage liability payments arising out of an occurrence.³ For purposes of statistical

² This description applies to occurrence-based liability policies, which are the type of policies used almost exclusively during the 1970s and early-to-mid 1980s. Since that time, insurers have also offered so-called “claims made” policies, which cover the insured for claims made during the policy period rather than negligent acts that later lead to claims as in the case of occurrence policies. Our data base applies to occurrence policies, but our analytical approach also could be applied to claims made policies.

analysis, the loss amount for any given occurrence is the sum of the loss and loss adjustment expense. This is appropriate because liability policies cover adjustment expenses (such as legal fees) as well as loss payments. In the discussion to follow, the term loss is understood to refer to the sum of losses and adjustment expenses. We use data only through 1986 because of structural changes in the ISO databases that occurred at that time that makes construction of a continuous database of consistent loss measurement difficult. This data set is quite extensive and hence is sufficient to contrast the implications of the methodologies outlined in this paper.

It is important to emphasize that the database consists of paid claim amounts, mostly for closed claims. Hence, we do not need to worry about the problem of modeling inaccurate loss reserves. The use of paid claim data is consistent with most of the loss reserving literature (e.g., Taylor 2000) and is the same approach adopted by the ISO (1994, 1998, 2002).

In the data set, the date (year and quarter) of the occurrence is given, and, for purposes of our analysis, occurrences are classified by accident year of origin. For each accident year, ISO classifies claims by payout lag by aggregating all payments for a given claim across time and assigning as the payment time the dollar weighted average payment date, defined as the loss dollar weighted average of the partial payment dates. For example, if two equal payments were made for a given occurrence, the weighted average payment date would be the midpoint of the two payment dates. For closed occurrences,⁴ payment amounts thus represent the total amount paid. For open claims, no payment date is provided. For open occurrences, the payment amount provided is the cumulative paid loss plus the outstanding reserve. The overall database consists

³ Aggregating bodily injury and property damage liability claims is appropriate for ratemaking purposes because general liability policies cover both bodily injury and property damage.

of 470,319 claims.

The time between the occurrence and the weighted average payment date defines the number of lags. Lag 1 means that the weighted average payment date falls within the year of origin, lag 2 means that the weighted average payment date is in the year following the year of origin, etc. Open claims are denoted as lag 0.

By considering losses at varying lag lengths, it is possible to model the tail of the payout process for liability insurance claims. Modeling losses by payout lag is important because losses settled at the longer lags tend to be larger than those settled soon after the end of the accident period. And, as we will demonstrate, the distributions also tend to be riskier for longer lags.

Both censored and uncensored data are included in the database. Uncensored data represent occurrences for which the total payments did not exceed the policy limit. Censored data are those occurrences where payments did exceed the policy limit. For censored data, the reported loss amount is equal to the policy limit so the total payment is the policy limit plus the reported loss adjustment expense.⁵ Because of the presence of censored data, the estimation was conducted using equation (9).

The numbers of occurrences by accident year and lag length are shown in Table 1. The number of occurrences ranges from 17,406 for accident year 1973 to 49,290 for accident year 1983. Thirty-four percent of the covered events for accident year 1973 were still unsettled 14 years later. Overall, about 28% of the claims during the period studied were censored.

Sample mean, standard deviation, skewness and kurtosis statistics for the products

⁴ Closed occurrences are those for which the insurer believes there will be no further loss payments. These are identified in the database as observations for which the loss reserve is zero.

⁵ Loss adjustment expenses traditionally have not been subject to the policy limit in liability insurance. More recently, some liability policies have capped both losses and adjustment expenses.

liability data are presented in Table 2.⁶ As expected, the sample means generally tend to increase with the lag length, indicating that larger claims tend to settle later. Exceptions to this pattern occur at some lag lengths, a result that may be due to sampling error since a few large occurrences can have a major effect on the sample statistics. Standard deviations also have a tendency to increase with lag length, indicating greater variation in total losses for claims settling later. Means and standard deviations also increase by accident year of origin. Sample skewnesses tend to fall between 2 and 125, revealing significant departures from symmetry. The sample kurtosis estimates indicate that these distributions have quite thick tails.

3.2 Estimated Severity Distributions

The products liability occurrence severity data are modeled using the generalized beta of the second kind, or *GB2*, family of probability distributions. Based on the authors' past experience with insurance data and an analysis of the empirical distribution functions of the products liability data, several members of the *GB2* family are selected as potential products liability severity models. We first discuss fitting separate distributions for each accident year and lag length of the runoff triangle. We then discuss and compare two alternative approaches for obtaining the overall severity of loss distribution – (1) the conventional approach, which involves fitting a single distribution to aggregate losses for each accident year; and (2) the use of the estimated distributions for the runoff triangle payment lag cells to construct a mixture distribution for each accident year. In each case, we propose the use of *discounted severity distributions* that reflect the time value of money. This is an extension of the traditional approach where loss severity distributions are fitted to loss data without recognizing the timing

of the loss payments, i.e., all losses for a given accident year were treated as undiscounted values regardless of when they were actually settled. The use of discounted severity distributions is consistent with the increasing emphasis on financial concepts in the actuarial literature, and such distributions could be used in applications such as the market valuation of liabilities.

3.2.1 Estimated Loss Distributions By Cell

In the application considered in this paper, separate distributions were fit to the data representing occurrences for each accident year and lag length. In other words, for 1973 the distributions were fit to occurrences settled in lag lengths 0 (open) and 1 (settled in the policy year) to 14 (settled in the thirteenth year after the policy year); for 1974, lag lengths 0 to 13, etc. The relative fits of the Weibull, the generalized gamma, the Burr 3, the Burr 12, and the four-parameter *GB2* are investigated. The parameter estimates for each distribution are obtained by the method of maximum likelihood. Convergence only presents a problem when estimating the generalized gamma.

The log-likelihood statistics confirm that the *GB2* provides the closest fit to the severity data for most years and lag lengths. This is to be expected because the *GB2*, with four parameters, is more general than the nested members of the family having two or three parameters. The likelihood ratio test helps to determine whether the *GB2* provides a statistically significant improvement over the other candidate distributions. Testing revealed that the Burr 12 provides a better fit than the generalized gamma and slightly better than the Burr 3. Thus, we chose the Burr 12 as the best three-parameter severity distribution. To provide an indication of its goodness of fit relative to two-parameter and four-parameter members of the *GB2* family,

⁶ It is to be emphasized that these are sample statistics corresponding to the data in nominal dollars. In many cases, the moments do not exist for the estimated probability distributions that are used to model the data.

Table 3 reports likelihood ratio tests comparing the Burr 12 and the Weibull distributions and the Burr 12 and *GB2* distributions.

The Weibull distribution is a limiting case of the Burr 12 distribution, as the parameter q grows indefinitely large. Likelihood ratio tests lead to the rejection of the observational equivalence of the Weibull and Burr 12 distributions at the one percent confidence level in 95 percent of the cases presented in Table 3. The null hypothesis in the Burr 12-*GB2* comparisons is that the parameter p of the *GB2* is equal to 1. The likelihood ratio tests lead to rejection of this hypothesis at the 5 percent level in 58 percent of the cases presented in Table 3 and also reject the hypothesis at the 1 percent level in 52 percent of the cases. Thus, the *GB2* generally provides a better model of severity.

The parameter estimates for the *GB2* are presented in the appendix. The mean exists if $\tilde{p} < 1/a < q$ and the variance if $\tilde{p} < 2/a < q$. The mean is defined in 96 out of the 118 cases presented in the appendix, but the variance is defined in only 33 cases. Thus, the distributions tend to be heavy-tailed. Further, policy limits must be imposed to compute conditional expected values in many cases and to compute variances in more than 70 percent of the cases.

Because the means do not exist for many cells of the payoff matrix, the median is used as a measure of central tendency for the estimated severity distributions. The medians for the fitted *GB2* distributions are presented in Table 4. As with the sample means, the medians tend to increase by lag length and over time. The median for lag length 1 begins at \$130 in 1973 and trends upward to \$389 by 1986. The lag 2 medians are almost twice as large as the lag 1 medians, ending at \$726 in 1985. The medians for the later settlement periods are considerably higher. For example, the lag 4 median begins at \$2,858 and ends in 1983 at \$8,038. Numbers along the diagonals in Table 4 represent occurrences settled in the same calendar year. For

example, 1986, lag 1 represents occurrences from the 1986 accident year settled during 1986; while 1985, lag 2 represents occurrences that were settled in 1986.

To provide examples of the shape and movement over time of the estimated distributions, GB2 density functions for lag 1 are presented in Figure 1. The lag 1 densities are shown for 1973, 1977 and 1981. The densities have the familiar skewed shape usually observed in insurance loss severity distributions. A noteworthy pattern is that the height of the curves at the mode tends to decrease over time and the tail becomes thicker. Thus, the probability of large claims has increased over the course of the sample period.

A similar pattern is observed in the density functions for claims settled in the second runoff year (lag 2), shown in Figure 2. Again, the tail becomes progressively heavier for more recent accident years. The density functions for lag 3 (Figure 3) begin to exhibit a different shape, with a less pronounced mode and much thicker tails. For the later lags, the curves tend to have a mode at zero or are virtually flat, with large values in the tail receiving nearly as much emphasis as the lower loss values. As with the curves shown in the figures, those at the longer lags have thicker tails. Thus, insurers faced especially heavy-tailed distributions when setting products liability prices in later years.

3.2.2 Estimated Aggregated Annual Loss Distributions

In the conventional methodology for fitting severity distributions, the entire data set for each accident year is used to estimate a single aggregate loss distribution. However, recall from the description of the data that claims for a given accident year are paid over many years subsequent to the inception of the accident year. To estimate a discounted severity distribution using the single aggregate loss distribution approach, we discount claims paid in lags 2 and higher back to the accident year. That is, the sample of claims used to estimate the distribution

for accident year j is defined as $y_{ijt}^d = y_{ijt} / (1+r)^{t-1}$, where y_{ijt} = the payment amount for claim i from accident year j settled at payout lag t , y_{ijt}^d = the discounted payment amount, and r = the discount rate. In the sample, $i = 1, 2, \dots, N_{jt}$ and $t = 1, \dots, T$, where N_{jt} = number of claims for accident year j settled at lag t , and T = number of settlement lags in the runoff triangle. Any economically justifiable discount rate could be used with this approach. In this paper, the methodology is illustrated using spot rates of interest from the U.S. Treasury security market.⁷

Based on the samples of discounted claims, maximum likelihood estimation provides parameter estimates of the GB2 distributions for each accident year. For example, of the 17,406 claims filed in 1973, 11,489 were closed during the next fourteen years and were discounted back from the time of closing to the policy year using spot rates of interest in force in the policy year. The 5,917 claims that remained open represent estimates of ultimate settlement amounts and were discounted back to the policy year assuming that they would settle in the year after the last lag year for the given policy year. Equation (9) was used for the parameter estimation, as there were claims that exceeded the policy limits and enter the estimation as censored observations. The appendix presents the GB2 parameter estimates.

The primary purpose for estimating the aggregated distributions is to compare it to the mixture distributions in the next section. Still, it is interesting to note that the GB2 distribution seems to provide the best model of the aggregated losses. We calculated likelihood ratio statistics for testing the hypotheses that the GB2 is observationally equivalent to the BR3 and BR12 distributions. In ten of the fourteen years, the GB2 provides a statistically significant

⁷ We extracted spot rates of interest from the Federal Reserve H-15 series of U.S. Treasury constant maturity yields. This data series is available at <http://www.federalreserve.gov/releases/h15/data.htm>. For each claim that was discounted we used the spot rate of interest as of the policy year with maturity equal to the lagged delay until the claim was settled.

improvement in fit relative to both the BR3 and BR12. In one year, 1978, the observed differences between the GB2 and neither the BR3 nor BR12 are statistically significant at conventional levels of significance. In 1973, 1976, and 1977, the GB2 provides a significant improvement relative to the BR3, but not relative to the BR12. Thus, the estimated GB2 distribution appears to generally provide a more accurate characterization of annual losses than any of its special cases.

3.2.3 Estimated Mixture Distribution for Aggregate Loss Distributions

A more general formulation of an aggregate loss distribution can be constructed using a mixture distribution. In this section, we present the mixture distribution for the undiscounted case, followed by the discounted mixed severity distribution. The section concludes with a brief discussion of our Monte Carlo simulation methodology.

The undiscounted mixture distribution can be developed by first assuming that each year in the payout tail may be modeled by a possibly different distribution. It will be assumed that the distributions come from a common family ($f(y_t; \theta_t)$), with possibly different parameter values (θ_t) where the subscript “ t ” denotes t^{th} cell in the payout tail and y_t is the random variable loss severity in cell t for a given accident year.⁸ The next step involves modeling the probability of a claim being settled in the t^{th} year, π_t , as a multinomial distribution. In the undiscounted case, the aggregate loss distribution is then obtained from the mixture distribution given by

$$f(y; \theta) = \sum_{t=1} \pi_t f(y_t; \theta_t) \quad (10)$$

Note that if θ_t is the same value for all cells, then the aggregate distribution would be the same as

⁸ To simplify the notation, accident year subscripts are suppressed in this section. However, the y_t are understood to apply to a particular accident year.

obtained by fitting an aggregate loss distribution $f(y; \theta)$ to the annual data. However, as mentioned, we find that the parameters differ significantly by cell within the payout triangle.

To obtain the discounted severity distribution for the mixture case, it is necessary to obtain the distributions of the discounted loss severity random variables, $y_t^d = y_t / (1+r)^{t-1}$. With the discount factor r treated as a constant, a straightforward application of the change of variable theorem reveals that discounting involves the replacement of the scale parameter of the GB2 distribution (equation (1)) by $b_t^d = b_t / (1+r)^{t-1}$, where b_t = the GB2 scale parameter for runoff cell t , and b_t^d = the scale parameter for the discounted distribution applicable to cell t .⁹

We estimate the parameters of the multinomial distribution, π_i , using the actual proportions of claims settled in each lag in our data set. The estimate of π_1 is the average of the proportions of claims actually settled in lag 1, the estimate of π_2 is the average of the proportions of claims actually settled in lag 2, etc. The estimate for π_0 is given by

$$\pi_0 = 1 - \sum_{i=1}^n \pi_i . \quad (11)$$

With the estimated cell severity distributions and the multinomial mixing distribution at hand, the mixture distribution was estimated using Monte Carlo simulation. The simulation was conducted by randomly drawing a lag from the multinomial distribution and then generating a random draw from the estimated severity distribution that corresponds to the accident year and lag. Each simulated claim thus generated is discounted back to the policy year in order to be

⁹ Treating the discount factor as a constant would be appropriate if insurers can eliminate interest rate risk by adopting hedging strategies such as duration matching and the use of derivatives. Cummins, Phillips, and Smith (2001) show that insurers use interest rate derivatives extensively in risk management. Modeling interest as a stochastic variable is beyond the scope of the present paper.

consistent with the data used in the estimated aggregate distribution.¹⁰ The estimated discounted mixture distribution is the empirical distribution generated by the 10,000 simulated claims for a given accident year, where each claim has been discounted to present value.

3.3 A Comparison of Estimated Aggregate Loss Distributions

As stated above, the aggregate loss distribution can be estimated in two ways. In this section, we compare the discounted mixture distribution based on equation (10) to a single discounted loss distribution fit to the present value of loss data aggregated over all lags. We will refer to these as the mixture distribution and the aggregated distribution, respectively. Both distributions were obtained using 10,000 simulated losses, from the mixture distribution and aggregated distribution, respectively.

With risk management in mind we illustrate the relationship between the mixture distribution and the aggregated distribution in Figure 4. Because the results are similar for the various accident years, we focus the comparison on the distributions for four accident years – 1973, 1977, 1981, and 1985. The figure focuses on the right tail of the distribution in order to illustrate the type of errors that could be made when failing to use the mixture specification of the aggregate loss distribution.

The most important conclusion based on Figure 4 is that the tails of the aggregated loss distributions are significantly heavier than the tails of the mixture distributions. Hence, the overall loss distribution applicable to the accident years shown in the figure appears to be riskier when the aggregate loss distribution is used than when the mixture distribution is used. We argue that the reason for this difference is that the aggregated distribution approach gives too

¹⁰ Equivalently, the discounted losses could be simulated directly from the GB2 distributions applicable to each payout cell, using the adjusted scale parameters b_t^d .

much weight to the large claims occurring at the longer lags than does the mixture distribution approach. In the aggregate approach, all large claims are treated as equally likely, whereas in the mixture approach the large claims at the longer lags are given lower weights because of the lower probability of occurrence of claims at the longer lags based on the multinomial distribution. In addition, rather than fitting the relatively homogeneous claims occurring within each cell of the runoff triangle, the model fitted in the aggregate approach is trying to capture both the relatively frequent small claims from the longer lags and the relatively frequent large claims from the longer lags, thus stretching the distribution and giving it a heavier tail. Thus, modeling the distributions by cell of the runoff triangle and then creating a mixture to represent the total severity distribution for an accident year is likely to be more accurate than fitting a single aggregate loss distribution (discounted or not) to the claims from an accident year regardless of the runoff cell in which they are closed. In our application, the aggregate approach overestimates the tail of the accident year severity distributions, but it is also possible that it would underestimate the tail in other applications with different patterns of closed claims. Of course, in many applications, such as reserving, it would be appropriate to work with the estimated distributions by payout cell rather than using the overall accident year distribution; but the overall discounted mixture distribution is also potentially useful in applications such as value-at-risk modeling based on market values of liabilities.

4. Summary and Conclusions

This paper estimates loss severity distributions in the payout cells of the loss runoff triangle and uses the estimated distributions to obtain a mixture severity distribution describing total claims from an accident year. We propose the use of discounted severity distributions, which would be more appropriate for financial applications than distributions that do not recognize the timing of

claims payments. We estimate severity of loss distributions for a sample of 476,107 products liability paid claims covering accident years 1973 through 1986. The claims consist of closed and open claims for occurrence based product liability policies. An innovation we introduce is to estimate distributions within each accident year/payment lag cell of the claims runoff triangle using a very general and flexible distribution, the generalized beta of the 2nd kind (GB2). Estimating distributions by cell is important because the magnitude and riskiness of liability loss distributions is a function both of the accident year of claim origin and the time lag between the occurrence of an event and the payment of the claim. Using a general severity distribution is important because conventional distributions such as the lognormal and gamma can significantly underestimate the tails of liability claims distributions.

The generalized beta family of distributions provides an excellent model for our products liability severity data. The estimated liability severity distributions have very thick tails. In fact, based on the GB2 distribution, the means of the distributions are defined for 81% of runoff triangle cells, and the variances are defined for only 28% of the cells. Thus, the imposition of policy limits is required in many cases to yield distributions with finite moments. The estimated severity distributions became more risky (heavy-tailed) during the sample period and the scale parameter for the early lags grew more rapidly than inflation. The results show that the gamma distribution, which has been adopted for theoretical modeling of claims by payout cell (e.g., Taylor 2000) would not be appropriate when dealing with the ISO products liability claims considered in this paper. Thus, it is appropriate to test the severity distributions that are to be used in any given application rather than making an assumption that the losses follow some conventional distribution.

Finally, we show that economically significant mistakes can be made if the payout tail is not accurately modeled. The mixture specification of the aggregate loss distribution leads to

significantly different estimates of tail probabilities than does the aggregated form of the aggregate loss distribution. In our application, the aggregate loss distribution tends to give too much weight to the relatively large claims from the longer lags and hence tends to overestimate the right tail of the accident year severity distribution. Such errors could create serious inaccuracies in applications such as dynamic financial analysis, reinsurance decision making, and other risk management decisions. Thus, the results imply that the mixture distribution and the distributions applicable to specific cells of the runoff triangle should be used in actuarial and financial analysis rather than the more conventional aggregate distribution approach.

Table 1: Products Liability Data Set — Numbers of Occurrences by Accident Year

Year	Payment Lag														OPEN	TOTAL
	1	2	3	4	5	6	7	8	9	10	11	12	13	14		
1973	5,011	2,876	743	528	439	395	318	151	144	201	215	195	207	66	5,917	17,406
1974	5,630	4,254	1,081	720	751	552	340	265	257	260	241	190	135		7,528	22,204
1975	8,124	4,919	1,128	1,003	805	661	419	320	259	491	235	201			7,782	26,347
1976	7,388	4,480	1,206	834	805	611	436	388	338	300	487				6,394	23,667
1977	7,616	5,023	1,359	965	796	661	607	424	286	228					4,756	22,721
1978	8,095	5,030	1,379	1,071	966	1,081	776	508	424						6,687	26,017
1979	9,529	8,658	2,196	1,939	1,525	1,281	857	961							9,516	36,462
1980	14,785	9,736	2,625	2,259	1,759	1,117	928								12,521	45,730
1981	16,439	10,599	2,850	2,191	1,629	1,333									11,596	46,637
1982	18,157	11,062	2,948	2,258	1,875										11,390	47,690
1983	19,198	12,870	3,360	2,512											11,350	49,290
1984	18,091	11,155	3,769												11,778	44,793
1985	16,671	10,525													9,876	37,072
1986	12,798														11,485	24,283

Table 3: Likelihood Ratio Tests

Burr 12 vs. Weibull														
Lag	1973	1974	1975	1976	1977	1978	1979	1980	1981	1982	1983	1984	1985	1986
0	766.00	1210.00	966.00	1286.00	214.00	366.00	1272.00	2340.00	1540.00	1060.00	460.00	320.00	540.00	1460.00
1	1560.00	1850.00	3250.00	2330.00	2204.00	2390.00	2640.00	3920.00	4140.00	4820.00	10700.00	4860.00	5280.00	6082.00
2	1168.00	1316.00	1672.00	1268.00	1516.00	1368.00	2930.00	2462.00	2594.00	3116.00	4820.00	3184.00	3526.00	
3	221.00	209.60	248.00	248.00	290.00	326.00	448.00	618.00	692.00	640.00	702.00	1102.00		
4	164.40	109.80	108.00	157.40	189.20	184.00	276.00	490.00	258.00	384.00	440.00			
5	57.00	46.60	85.60	124.40	131.40	70.00	196.00	180.00	370.00	294.00				
6	86.60	39.80	34.00	54.20	45.20	2.00	112.00	228.00	164.00					
7	45.00	21.40	12.40	62.20	2.80	22.60	1076.00	56.00						
8	12.00	18.00	13.80	21.00	24.00	24.40	68.00							
9	26.00	35.40	10.40	38.40	5.00	19.20								
10	3.40	27.20	111.80	2.80	0.40									
11	11.60	11.00	7.60	349.20										
12	38.40	67.40	186.40											
13	191.40	89.40												
14	40.82													
GB2 vs. Burr 12														
Lag	1973	1974	1975	1976	1977	1978	1979	1980	1981	1982	1983	1984	1985	1986
0	168.00	108.00	52.00	372.00	0.00	16.00	22.00	0.00	80.00	240.00	280.00	60.00	0.00	120.00
1	184.00	76318.00	262.00	246.00	244.00	210.00	226.00	400.00	520.00	440.00	260.00	420.00	280.00	190.00
2	54.00	126.00	114.00	98.00	102.00	134.00	160.00	140.00	64.00	190.00	60.00	140.00	114.00	
3	0.20	12.00	4.00	10.00	8.00	4.00	22.00	12.00	6.00	16.00	10.00	0.00		
4	0.20	3.20	0.20	0.40	0.40	0.00	2.00	2.00	2.00	10.00	0.00			
5	0.80	3.40	21.40	0.00	0.00	18.00	0.00	2.00	0.00	6.00				
6	10.80	10.60	12.20	5.00	0.80	48.00	0.00	0.00	0.00					
7	0.40	0.20	1.00	0.20	9.80	2.00	46.00	10.00						
8	5.40	3.40	0.40	6.60	1.00	2.00	2.00							
9	0.80	3.80	0.00	12.20	0.40	1.40								
10	0.00	0.20	72.20	2.60	2.60									
11	0.20	0.20	2.40	36.40										
12	1.60	1.20	8.20											
13	11.60	8.60												
14	0.00													

The reported likelihood ratio statistics are twice the difference in the log-likelihood function values. The likelihood ratio statistic in the reported tests has a chi-square distribution with one degree of freedom. The hypothesis being tested in the comparison of the Weibull and the Burr 12 is that q in the Burr 12 distribution is infinite. The hypothesis in the comparison of the GB2 and the Burr 12 is that $p = 1$ in the GB2. A value larger than 3.8 represents rejection of the hypothesis at the 95% confidence level and values larger than 6.5 represent rejection of the hypothesis at the 99% confidence level. McDonald and Xu (1992) note that rejection of a hypothesis of an infinite parameter value at the 95% confidence level using the traditional likelihood ratio test is approximately equal to rejecting at the 98% confidence level with an appropriately modified test statistic.

Figure 1
GB2 Lag 1 Density Functions




Figure 2
GB2 Lag 2 Density Functions




Figure 3
GB2 Lag 3 Density Functions








Figure 4
Aggregate GB2 Loss Distributions Using a
Mixture Form and an Aggregated Form



Note: The solid line represents the cumulative loss distribution estimated by using a multinomial distribution to simulate losses drawn from the individual years of the payout tail. The dashed line represents an estimate of the cumulative loss distribution where all years of the tail are discounted to the time of pricing and a single GB2 distribution is fit.

Figure 4 – Continued



Note: The solid line represents the cumulative loss distribution estimated by using a multinomial distribution to simulate losses drawn from the individual years of the payout tail. The dashed line represents an estimate of the cumulative loss distribution where all years of the tail are discounted to the time of pricing and a single GB2 distribution is fit.

References

- Aiuppa, Thomas A. (1988). "Evaluation of Pearson curves as an approximation of the maximum probably annual aggregate loss." *Journal of Risk and Insurance* 55: 425-441.
- Bookstaber, Richard M. and James B. McDonald (1987). "A General Distribution for Describing Security Price Returns," *Journal of Business* 60: 401-424.
- Butler, Richard J. and James B. McDonald, (1987). "Some Generalized Mixture Distributions with an Application to Unemployment Duration," *Review of Economics and Statistics*, 69: 232-240.
- Cummins, J. David, Georges Dionne, James B. McDonald, and B. Michael Pritchett (1990). "Applications of the GB2 Family of Distributions In Modeling Insurance Loss Processes." *Insurance: Mathematics and Economics* 9: 257-272.
- Cummins, J. David, Christopher M. Lewis, and Richard D. Phillips (1999). "Pricing Excess of Loss Reinsurance Contracts Against Catastrophic Loss," in Kenneth Froot, ed., *The Financing of Catastrophe Risk* (Chicago: University of Chicago Press).
- Cummins, J. David, Richard D. Phillips, and Stephen D. Smith (2001). "Derivatives and Corporate Risk Management: Participation and Volume Decisions in the Insurance Industry." *Journal of Risk and Insurance* 68: 51-91.
- Elderton, W.P. and N.L. Johnson (1969). *Systems of Frequency Curves*. Cambridge University Press,
- Hogg, Robert and Stuart Klugman (1984). *Loss Distributions*. New York: John Wiley & Sons.
- Insurance Services Office (1994). "Products/Completed Operations Liability Indicated Mixed Pareto Increased Limits Factors," New York.
- Insurance Services Office (1998). "Products/Completed Operations Liability Increased Limits Data and Analysis," New York.
- Insurance Services Office (2002). "Products/Completed Operations Liability Increased Limits Data and Analysis," New York.
- Johnson, Norman L. and Samuel Kotz (1970). *Continuous Univariate Distributions*, Vol 1. Wiley New York.
- Klugman, Stuart A., Harry H. Panjer, and Gordon E. Willmot (1998). *Loss Models: From Data to Decisions* (New York: John Wiley & Sons).
- Mack, Thomas (1991). "A Simple Parametric Model for Rating Automobile Insurance or Estimating IBNR Claims Reserves," *Astin Bulletin* 21: 93-109.

- McDonald, James B. (1984). "Some Generalized Functions of the Size Distribution of Income." *Econometrica* 52: 647-663.
- McDonald, James B. and Yexiao J. Xu (1995). "A Generalization of the Beta of the First and Second Kind." *Journal of Econometrics*, 66:133-152.
- Paulson, A.S. and N.J. Faris (1985). "A Practical Approach to Measuring the Distribution of Total Annual Claims." In J.D. Cummins, ed., *Strategic Planning and Modeling in Property-Liability Insurance*. Norwell, MA: Kluwer Academic Publishers.
- Ramlau-Hansen, Henrik (1988). "A Solvency Study in Non-life Insurance. Part 1. Analysis of Fire, Windstorm, and Glass Claims." *Scandinavian Actuarial Journal*, pp. 3-34.
- Reid, D.H. (1978). "Claim Reserves in General Insurance," *Journal of the Institute of Actuaries* 105: 211-296.
- Taylor, Gregory C. (1985). *Claims Reserving in Non-Life Insurance* (Amsterdam: North-Holland).
- Taylor, Gregory C. (2000). *Loss Reserving: An Actuarial Perspective* (Boston: Kluwer Academic Publishers).
- Ventor, Gary C. (1984). "Transformed beta and gamma functions and aggregate losses." Reprinted from: *Proceedings of the Casualty Actuarial Society*, Vol 71. Recording and Statistical Corporation, Boston, MA.
- Wiser, Ronald F., Jo Ellen Cockley, and Andrea Gardner (2001). "Loss Reserving," in *Foundations of Casualty Actuarial Science*, 4th ed. (Arlington, VA).
- Wright, T.S. (1990). "A Stochastic Method for Claims Reserving in General Insurance," *Journal of the Institute of Actuaries* 117: 677-731.