Uncertainty in Hurricane Risk Modeling and Implications for Securitization

David Miller, Ph.D.
Vice President, Guy Carpenter & Company

Abstract

This paper presents a simple procedure for quantifying the effect of the historically limited meteorological record on insurance losses calculated by hurricane catastrophe simulation models. Using a standard actuarial approach, the uncertainty in the historical record can be decomposed into uncertainty in the mean annual hurricane frequency, and uncertainty in the severity, where the severity corresponds to the distribution of insurance losses resulting from a single random hurricane event. Uncertainty in the mean annual frequency can be estimated directly from the historical data. Uncertainty in the severity is more difficult to evaluate, but can be approximated using a parametric bootstrap method which captures the effects of the finite historical record. Confidence intervals on the insurance loss as a function of return period are calculated for several different portfolios. These results are placed in the context of a recent insurance-linked securitization transaction, and are found to be consistent with other independent measurements of pricing uncertainties in similar instruments. The historical volatility of one-year corporate bond default rates is also compared to the magnitude of the uncertainty in the hurricane simulation models. This comparison suggests that volatility of one-year corporate bond default rates is greater than uncertainty in the hurricane models.

I. Introduction

Hurricane catastrophe simulation (CAT) models represent an attempt to apply a climatological pattern of hurricane activity to an existing insurance inventory.[Friedman,1984] These models typically assume a stationary climate consistent with the average historical record as the basis for stochastic hurricane simulation. Because hurricane records in the Atlantic basin span only the past 100 years, any simulation of hurricane activity in this region will ultimately be limited by this finite number of historically observed hurricane events. In particular, the sampling variability associated with developing the meteorological hazard component of a CAT model from the finite number of historical hurricane observations represents a fundamental, epistemological source of uncertainty in the estimation of insurance losses.

The inability to predict an outcome can be decomposed into two sources: Randomness and uncertainty. Randomness, or process risk, is due to the fluctuating nature of the modeled phenomenon, and may be accurately accounted for by stochastic simulation models. Uncertainty, on the other hand, is due to limited modeling knowledge and data, and in general is not implicitly accounted for by stochastic simulation models. For example, the randomness associated with each roll of a pair of dice is conceptually distinct from the uncertainty associated with not knowing whether the dice are 'fair'. While this distinction can be difficult to make in some situations, conflating uncertainty with randomness in stochastic simulation models may have serious implications for bias.[Major,1998]

The effect of uncertainty within a given model can further be specified in terms of model specification error, reflecting uncertainty in the results due to the potentially incorrect choice of probability distribution family, and parameter risk, associated with the incorrect choice of parameters given the correct family. The goal of this paper is to quantify the uncertainty associated with the historically limited meteorological record. While there are many sources of uncertainty in CAT models, the finite number of historical hurricane observations underlies sources of both model specification error as well as of parameter risk. Analysis of the effects of the finite historical record on estimating insurance losses will therefore provide a lower bound estimate for uncertainty in hurricane CAT models.

Using a standard actuarial approach, the uncertainty in the historical record can be decomposed into uncertainty in the mean annual hurricane frequency, and uncertainty in the severity, where the severity corresponds to the distribution of insurance losses resulting from a single (random) hypothetical hurricane event. Decomposition into frequency and severity components is made under the assumption that the severity is independent of frequency. Uncertainty in the mean annual frequency is estimated directly from the historical data. Uncertainty in the severity is more difficult to evaluate, but can be approximated using a parametric bootstrap method which captures the effects of the finite historical record. Using this technique, confidence intervals on the exceedance probability distributions produced by hurricane CAT simulation models are calculated for three different portfolios. Although a particular CAT model is used in part of this analysis, this uncertainty analysis technique may be applied to any CAT model. Regardless of computational details, all CAT models rely on the 100 year historical record and are therefore subject to the uncertainty resulting from this limited dataset.

There has been much recent attention given to insurance-linked securitizations, or so-called 'catastrophe bonds.'[Buckley and Orr, 1998] These instruments provide a mechanism for catastrophe risk transfer as an alternative to traditional reinsurance. In order to place the analysis of uncertainty in hurricane catastrophe simulation models within the context of these instruments, three comparisons are made. First, the procedure for evaluating confidence intervals on the exceedance probability distributions developed here is placed in the context of rating CAT bonds. Second, the results of this analysis are compared with an independent analysis of uncertainties in pricing CAT securities.[Moore, 1998] Third, the magnitude of the volatility in historical one-year corporate bond default rates is compared to an equivalent measure of uncertainty in CAT bond default rates. This comparison is motivated by a desire to gain an understanding of how investors might compare the new CAT bonds with more traditional investment vehicles.

II. Hurricane Catastrophe Simulation Models

Approximately[1] 210 hurricanes have made landfall in the eastern US during the past 100 years, resulting in a mean annual frequency of 2.1 landfalling hurricanes per year. The relative paucity of hurricane events, along with the complications brought about by rapidly changing demographics, are the primary reasons that catastrophe simulation models are required for accurate forecasting of hurricane losses. Traditional trend-and-develop approaches for extrapolating the probabilities of insurance losses resulting from hurricanes are replaced by simulation models, which attempt to separately model the underlying meteorology.

Hurricane simulation models have three major components: A hazard component, an engineering component and an insurance component. The hazard component is used to generate the distribution of hurricane wind speeds at each geographical site in the portfolio that are produced by a particular storm or set of storms. The engineering component takes these wind fields, along with site specific information pertaining to all structures in the insurance inventory, and calculates damage rates at each physical location in the portfolio and for each structure and coverage type within the policy. The insurance component applies deductibles and limits to the site specific damages calculated by the engineering component. Policy specific deductibles and limits, along with all remaining insurance structure, are then applied to the realized losses, and aggregated over all policies for each simulated event.

The simulation model can be run in either a scenario or in a probabilistic mode. In the scenario mode, the hurricane wind field resulting from a single hurricane event is simulated by the hazard component. The simulated event may be based on an historically-observed hurricane event, or may be some hypothetical event. In the probabilistic mode, the hazard module simulates wind fields corresponding to the climatological ensemble of possible hurricane events, of which the historical events are considered a random sample.

[1] This lack of precision is a function of definitional questions surrounding the term 'landfall.'

The output of the probabilistic mode of the hurricane CAT simulation model can be represented by the per-occurrence exceedance probability curve (EP) for a given portfolio. This represents the exceedance distribution of the maximal loss occurrence in the course of the year for the given portfolio, or

$$EP\ (l) = P\{\text{A loss greater than } l \text{ will occur during the year}\}. \qquad (1)$$

Using a standard actuarial approach, the distribution of losses which are simulated by the hurricane CAT model and which are represented by the EP curve can be decomposed into separate frequency and severity components. [Hogg, 1984] The frequency corresponds to the distribution of the number of hurricanes occurring within a one year period. The severity corresponds to the distribution of losses conditional on the occurrence of a single, random hurricane. The cumulative distribution function (CDF) corresponding to the severity is defined as

$$F_S\ (l) = P\{\text{Given that a hurricane occurs, the loss will not exceed } l\ \}. \qquad (2)$$

Typically, the Poisson distribution is used to model event frequency. If the mean frequency is λ, and frequency is independent of severity, then the number of events occurring within the course of a year with loss greater than an amount l is also Poisson with mean equal to $\lambda\ (1 - F_S\ (l)\ )$. The per-occurrence exceedance probability can then be written as

$$EP(l) = 1 - \text{Exp}[-\lambda\ (1 - F_S\ (l)\ )]. \qquad (3)$$

The EP curve represents the effects of the climatological ensemble of hurricanes on a given portfolio. The ensemble, of course, is not known, but rather is estimated from the historically observed hurricanes. If a single hurricane event can be specified by a vector $\boldsymbol{\theta}$ of storm parameters, then the set of historically observed hurricane events can be represented as $\Theta_{Historical} = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \ldots \boldsymbol{\theta}_N)$, where N is the number of hurricane observations which have occurred over the past 100 years. Each vector $\boldsymbol{\theta}_i$ contains the complete list of parameters used to describe the physical properties of the i^{th} historical hurricane, for example, landfall location, initial pressure drop, translation speed, etc.. The dependence of the exceedance probability on the historical storms can then be written as

$$EP(l\ , \boldsymbol{A} \mid \Theta_{Historical}) = 1 - \text{Exp}[-\lambda(\Theta_{Historical})\ \ (1 - F_S\ (l\ , \boldsymbol{A} \mid \Theta_{Historical}\ )\ )], \qquad (4)$$

where the dependence of the EP on the particular choice of portfolio $\boldsymbol{A}$ is also included. The matrix $\boldsymbol{A}$ contains all relevant information on the particular portfolio of interest, including site locations, structural information, total insured value as well as all relevant insurance information.

The mean annual frequency $\lambda(\Theta_{Historical})$ can simply be calculated as the total number of hurricanes in the historical record, N, divided by the number of years in the historical record, n. In contrast, the exact dependence of $F_S\ (l\ , \boldsymbol{A} \mid \Theta_{Historical})$ on $\Theta_{Historical}$ is considerably more complicated. In general, F_S is constructed in such a way as to reproduce the meteorological variability of the observed historical hurricane events while at the same time minimizing the sampling variance associated with the finite number of historical hurricane events. For example, a typical hurricane CAT model will not simulate just the 210 or so historically occurring hurricane events, but will rather simulate a much larger number of hurricane events, providing a more smoothly varying geographical and intensity coverage. The implementation of this procedure is highly model dependent, and forms the ‘art’ of CAT modeling.

III. Estimating confidence intervals on the exceedance probability

The EP curve produced by the CAT simulation model is generated, ultimately, from the historical hurricane parameters as specified by $\Theta_{Historical}$. Implicit in $F_S\ (l\ , \boldsymbol{A} \mid \Theta_{Historical}\ )$ is the inclusion of randomness related to the storm parameters as well as to vulnerability (engineering) response. In general, it will not include the effects of model specification uncertainty. It may or may not reflect parameter uncertainty, depending on the modeler's 'art.' We will assume here that EP incorporates only randomness, and not uncertainty.

From Eq. 4, it can be seen that uncertainty in the EP curve is related to uncertainty in the mean annual hurricane frequency compounded with uncertainty in the severity component. In this section, uncertainty in the frequency and severity components resulting from the finite historical record is estimated. Uncertainty in the EP curve, in the form of confidence intervals on the base-case EP curve, is then calculated from the frequency and severity uncertainties. The base-case refers to the EP curve defined by Eq. 3 using the historical mean annual hurricane frequency λ and the empirically-driven severity distribution $F_S(l, \boldsymbol{A} \mid \Theta_{Historical})$ that is simulated by the CAT model for a given portfolio $\boldsymbol{A}$. In the bootstrap literature,[Efron, 1993] the base-case is referred to as the plug-in EP curve, denoted $\hat{EP}(l)$.

Uncertainty in the mean annual hurricane frequency is represented by the standard error of the mean, which is determined directly from the historical data as

$$\sigma_{\bar{x}} = \sqrt{\frac{1}{n(n-1)} \sum_{i=1}^{n} (x_i - \bar{x})^2}, \tag{5}$$

where n is the number of years in the historical record, and x_i is the number of hurricanes occurring in year i. It is assumed that the hurricane frequency is time independent, i.e., that the $\{x_i\}$are independent and identically distributed (iid). This assumption is independently confirmed. For n = 96, $\bar{x}$ = 2.07 and $\sigma_{\bar{x}}$ = 0.159. This is in agreement with the theoretical standard error (se) on the mean calculated by assuming Poisson statistics with mean frequency $\lambda = \bar{x}$, yielding an se = $(\lambda/n)^{1/2}$ = 0.147. A t-distribution is used to model the sampling distribution of mean frequencies. Figure 1 shows the results of a bootstrap simulation of the historical data $\{x_i\}$ with 500 replications, along with the t-distribution calculated with the empirical value for $\sigma_{\bar{x}}$. The agreement between the t-distribution and the bootstrap simulation is good.

The bootstrap technique [Efron, 1993] can be used to estimate uncertainty in the severity distribution due to the finite number of historical hurricane observations. The dependence of the severity distribution on the historical record is represented as $F_S(l, \boldsymbol{A} \mid \Theta_{Historical})$. The bootstrap technique makes the assumption that the observed dataset, namely $\Theta_{Historical}$, is a representative subset of potential outcomes from some underlying distribution, so that random subsamples from the observed dataset, namely $\Theta_{Historical}$, are themselves representative subsets of potential outcomes. Assuming a stationary climate, the underlying distribution can be thought of as the distribution of hurricane parameters representative of an infinitely long hypothetical historical record. Each bootstrap replication represents an equivalent realization of the historical record, and consists of random draw, with replacement, of N hurricanes from the observed record of $\Theta_{Historical} = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \ldots \boldsymbol{\theta}_N)$. Each bootstrap replication can be represented as $\Theta^{*b} = (\boldsymbol{\theta}_1^{*b}, \boldsymbol{\theta}_2^{*b} \ldots \boldsymbol{\theta}_N^{*b})$ with b = 1...B and B = 500. Specifically, each $\boldsymbol{\theta}_i^{*b}$ is drawn randomly from the collection $\Theta_{Historical} = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \ldots \boldsymbol{\theta}_N)$. Confidence intervals on the severity can then be determined from the B bootstrap replication of $F_S^{*b}(l, \boldsymbol{A} \mid \Theta^{*b})$. For example, at each value of the loss *l* there are B bootstrap values of the severity, from which percentiles on F_S can easily be created by a simple sorting of the values. The resulting confidence intervals on the severity reflect the uncertainty in the EP curve due to the limited number of historical events.

Unfortunately, this procedure is prohibitively time consuming. In particular, the need to create $F_S^{*b}(l, \boldsymbol{A} \mid \Theta^{*b})$ for each bootstrap replication, depending as it does on the CAT model-building 'art,' makes this approach impractical. Essentially, each $F_S^{*b}(l, \boldsymbol{A} \mid \Theta^{*b})$ represents a new CAT simulation model, created as if the observed historical event set had been Θ^{*b} instead of $\Theta_{Historical}$ – surely an involved process. Therefore, an alternate approach is taken, which utilizes the parametric bootstrap technique. In this approach, a severity distribution $F_S(l, \boldsymbol{A} \mid \Theta_{Historical})$ is calculated for a given portfolio using a particular choice of CAT model. This empirical severity distribution $F_{SE}(l)$ is then fit to a (simpler) parametric model, with the best-fit to the parametric model defined as $F_{SP}(l)$. The parametric model is chosen based on its ability to fit the empirical severity distribution. Once $F_{SP}(l)$ is specified, a Monte Carlo algorithm is used to simulate 210 points from this parametric distribution. The number of points used in the simulation corresponds to the number of historical hurricanes N. B = 500 Monte Carlo replications of 210 points each are created in the analysis. For each collection of 210 points, the parametric model is used to find a new best-fit, resulting in a new parametric severity distribution $F_{SP}{*}^{b} (l)$ for each replication b. The resulting parametric severity distributions are approximations to the $F_S{*}^{b} (l)$, and confidence intervals on the severity can be created as described above.

The parametric bootstrap calculation provides a reasonable approximation to the ideal bootstrap calculation described above. The simulation of 210 severity points in each of the parametric bootstrap replications is analogous to bootstrapping the historical storm elements. The advantage of the parametric technique is the reduction in computational load resulting from elimination of the need to create a new severity $F_S{*}^{b} (l , \boldsymbol{A} \mid \Theta{*}^{b})$ (a new CAT model!) for each bootstrap replication.

Confidence intervals on the severity are calculated for three different portfolios $\boldsymbol{A}$ using the parametric bootstrap technique. A commercially available hurricane CAT simulation model is used to simulate losses for each portfolio. The resulting empirical severity distributions $F_{SE}(l)$ simulated by the CAT models for the three different portfolios are displayed in Fig. 2 as the thick solid lines. The portfolios correspond to (a) personal, (b) commercial and (c) specialty lines exposure, respectively. Note that all insurance portfolios are geographically distributed over a large area of the US.

The four parameter transformed beta distribution, with an additional zero-mass term, is used to fit the empirical severity distributions as well as each of the parametric bootstrap replications. This parametric model is given by

$$F_{SP}{}^{TB}(x) = (1 - w) + w\,\beta\left(\frac{x^{\tau}}{\lambda + x^{\tau}}, \kappa, \alpha\right) \qquad (6)$$

where $F_{SP}{}^{TB}(x)$ is the parametric severity CDF described by the transformed beta function with zero mass term. The quantity (1 – w) corresponds to the zero mass, or the probability of having zero loss to the portfolio given the occurrence of a hurricane event, and β is the incomplete Beta function. The best-fit to the empirical severity distributions, $F_{SP}(l)$, is displayed as the respective thick dashed lines in Fig. 2. The thin dotted lines in Fig. 2 correspond to the parametric bootstrap replication severity distributions $F_{SP}{*}^{b} (l)$ plotted for every 10th bootstrap replication. A total of B = 500 replications are made for each portfolio. For all portfolios, the transformed beta distribution was able to fit the empirical severities to a 1% significance, as measured by the Kolmogorov-Smirnov statistic.

Plug-in (base case) EP curves $\hat{EP}(l)$ for the three portfolios are shown in Fig. 3 as the middle solid lines, where the exceedance probability EP has been transformed into a return period, with the return period T = 1/EP. In this figure the EP curve is represented by a log-log plot of loss in dollars versus return period in years. For example, an annual exceedance probability of 1% corresponds to a 100 year return period.

Confidence intervals on the plug-in $\hat{EP}(l)$ curve resulting from compounding uncertainty in the mean annual frequency and severity can be determined in the following way: First, tabulate evenly spaced percentiles on the mean annual frequency distribution, $\lambda^{(p)}$, where for example p = 5%, 10%, ... 95%. Next, at each value of the loss l, calculate the percentile severity curves $F_S{}^{(p')} (l)$ at p' = 5%, 10%, ... 95% from the B = 500 sorted values of $F_{SP}{*}^{b}(l)$ as depicted in Fig. 2. Now the argument to the exponential in Eq. 3 is just the product of λ and $(1 – F_S (l))$, so combining the 19 mean annual frequency percentiles $\lambda^{(p)}$ with the 19 percentile severity curves $F_S{}^{(p')} (l)$ results in $361 = 19^2$ curves corresponding to the product $\lambda^{(p)} (1 – F_S{}^{(p')} (l))$. Said another way, there are 361 values of the exceedance probability $EP^k(l)$ at each value of the loss l, where k = 1...361. Because the resulting 361 curves for $EP^k(l)$ may cross, confidence intervals on the

EP curve must be determined from the sorted values of the $EP^k(l)$ for each value of the loss l. In this way, percentile exceedance probability curves $EP^{(p)}(l)$ can be formed.

Percentile EP curves are shown in Fig. 3 as a log-log plot of loss versus return period along with the plug-in EP curves. The upper solid lines corresponds to the 95th percentile curve $EP^{(.95)}(l)$ and the lower solid lines corresponds to the 5th percentile curve $EP^{(.05)}(l)$. The 90% confidence intervals for the EP curve determined in this way are based on the uncertainty in the mean annual frequency and in the severity resulting from the limited historical dataset. The dashed lines in Fig. 3 correspond to the case where only uncertainty in the mean annual frequency is used to determine a 90% confidence interval for the EP curve. Calculating confidence intervals on the EP curve based only on uncertainty in the severity produces curves virtually identical to the curves corresponding to the inclusion of both frequency and severity uncertainty. This result indicates that uncertainty in the mean annual frequency is relatively insignificant compared to uncertainty in the severity.

The confidence interval shown in Fig. 3 can be represented in another way by normalizing the percentile EP curves to the plug-in EP curve. In Fig. 4, the ratio $L^{(p)}/\hat{L}$ versus return period is plotted for the three portfolios (a) through (c). The quantity $L^{(p)}$ is the loss percentile for percentile p, and $\hat{L}$ is the plug-in estimate for the loss at the respective return period, obtained from the plug-in EP curve $\hat{EP}(l)$. The ratios which are greater than 1 correspond to the percentile p = 95% and the ratios which are less than 1 correspond to p = 5%. The solid, dashed and dotted lines correspond to portfolios (a), (b) and (c), respectively. The area between the two curves for each portfolio represents the 90% confidence interval on the normalized loss as a function of return period.

There are several interesting features in this plot. At very low return periods (< 10 years) the 95th percentile ratio diverges and the 5th percentile ratio goes to zero. This is due to the fact that $\hat{L}$, the plug-in estimate for the loss, is going to zero, while the numerators in the ratio $L^{(p)}/\hat{L}$ are going to zero at a slower or faster rate than the denominator, respectively, for p = 95% and p = 5%. In addition, for all curves, there appears to be some feature near a return period of ~ 10 years. Perhaps most interestingly, all three of the ratios are approximately constant for return periods greater than ~ 80 years. For all portfolios, the normalized 5th and 95th percentile loss curves are approximately 0.5 and 2.5 times the base case, respectively, between return periods of 80 to 1000 years.

There is a simple argument that can be used to support the magnitude of the confidence intervals observed in Fig. 4. Imagine trying to evaluate the damage to a single site somewhere along the US coast. In the past 100 years, there have been approximately 210 hurricanes that have made landfall along the entire coastline. Consider that only those historical hurricane tracks passing within a window of about 200 nautical miles (nmi) have any chance of causing damage at any given site. With a total coastal length of approximately 3100 nmi (from the Mexico-US border to the US-Canada border), on average 14 of those 210 storms would have passed within a 200 nautical mile window of this site. The average number of storms affecting the site is therefore $\mu = 0.14$ per year. Using Poisson statistics to determine the standard error (se) of the number of events affecting the site leads to se = $(0.14/100)^{1/2} = 0.04$. The resulting coefficient of variation CoV = se/μ = 30%. Assuming a normal distribution for the number of events affecting the site, and assuming that the uncertainty in the severity is due solely to the uncertainty in the number of events affecting the site, then the 90 percent confidence interval for the loss is given by ± (1.65)(30%) = ± 50%. This is consistent with results of Fig. 4. Additional sources of uncertainty will tend to increase the width of this 90 percent confidence interval.

IV. Implications for uncertainty in CAT bond pricing

The implications of uncertainty in CAT simulation models on the rating and pricing of CAT bonds are explored here by placing the results of the previous section within the context of a recent CAT bond rating. The specific case to be examined is the United Services Automobile Association (USAA) / Residential Re contract, which provided 80% coverage of $500m in losses in excess of USAA's $1b retention. The bond side of this transaction was rated by Duff & Phelps Credit Rating Company (DCR) with Applied Insurance Research (AIR) as the model consultant.[Buckley and Orr, 1998]

DCR rates CAT bonds based on the values of three statistics: probability of first dollar of loss (FL), probability of depletion (last dollar of) loss (DL), and the expected loss (EL), where EL is normalized to the layer size. The quantities $FL_\$$, $DL_\$$ and $EL_\$$ represent the actual dollar values of the lower limit, upper limit and expected loss of the CAT bond layer, respectively. The normalized expected loss can be expressed as $EL = EL_\$ /(DL_\$ - FL_\$)$. FL and DL are the exceedance probabilities associated with respective loss to the portfolio. Assuming that the EP curve is available from a CAT model simulation of the portfolio losses, then $FL = EP(FL_\$)$ and $DL = EP(DL_\$)$. EL can be calculated using a standard actuarial expression, as

$$\mathrm{EL} = (\mathrm{DL}_\$ - \mathrm{FL}_\$)^{-1} \times \int_{\mathrm{FL}_\$}^{\mathrm{DL}_\$} (x - \mathrm{FL}_\$)\, f(x)dx + \int_{\mathrm{DL}_\$}^{\infty} f(x)dx\,. \tag{7}$$

The function *f(x)* is the probability density function (pdf) corresponding to the exceedance probability curve EP(x) determined, in this case, from AIR's hurricane CAT simulation model. Using the EP curve determined by AIR's CAT model analysis of the USAA portfolio and taking $FL_\$ = \$1b$ and $DL_\$ = \$1.5b$ leads to FL = 1.35%, DL = 0.55% and EL = 0.72%.

The three portfolios (a) - (c) used in the EP confidence level analysis in this paper clearly have different exposures than the USAA portfolio rated by DCR both in total insured value (TIV) as well as in geographical distribution. However, the exposure in each of these three portfolios can be qualitatively 'normalized' to the USAA portfolio by calculating an $FL_\$$ and $DL_\$$ for each of these portfolios based on the respective portfolio's EP curve and the USAA FL = 1.35% and DL = 0.55% values. Values of EL calculated for each of the portfolios (a) - (c) using Eq. 7 should then be comparable to the value of EL calculated for the USAA transaction. This comparison is shown in Table 1. Results of the USAA/Res Re analysis are given in the first row of the table; the corresponding results for portfolios (a) - (c) are given below that. Values for $FL_\$$ and $DL_\$$ in the table are in billions of dollars; values for EL are in percent. Values for EL are denoted by $\hat{EL}$ since they are calculated from the respective plug-in (base case) EP curve. As can be seen from the table, the values of $FL_\$$ and $DL_\$$ for portfolio (a) are similar in size to the corresponding values for the USAA portfolio, indicating that these two portfolios probably contain similar TIV. In contrast, values of $FL_\$$ and $DL_\$$ for portfolio (c) are much smaller. Despite these differences in the values of $FL_\$$ and $DL_\$$, values of EL for portfolios (a) – (c) are very similar to the USAA portfolio result.

Values of EL calculated in Table 1 are based on the respective portfolio's plug-in EP curve, and therefore reflect randomness associated with the meteorological variability and the vulnerability response. However, these values of EL do not include the effects of model specification or parameter uncertainty. Confidence intervals on EL can be calculated from the parametric bootstrap analysis described in the previous section. Specifically, the B = 500 bootstrap replications of the severity distributions $F_S{*}^b(l)$ created for each of the portfolios (a) - (c) are combined with the 19 mean annual frequency percentiles $\lambda^{(p)}$ resulting in 9,500 = (19)(500) equiprobable EPcurves. This collection of EP curves is used to determine confidence intervals on EL. For each portfolio, each of the 9,500 bootstrap EP curves is numerically integrated over the range specified by the respective values of $FL_\$$ and $DL_\$$. The 99th percentile EL value is then calculated from the sorted results and can be expressed as a difference, normalized to the plug-in EL value as

$$\Delta^{(.99)} = \frac{EL^{(.99)} - \hat{EL}}{\hat{EL}}. \qquad (8)$$

The statistic $\Delta^{(.99)}$ is a measure of pricing uncertainty, or more specifically, uncertainty of the expected loss calculated for a particular CAT bond transaction. This statistic is closely related to uncertainty in CAT bond default rates. Default rates for CAT bonds may be defined as the attachment probability, or probability of first loss (FL), and uncertainty of FL is similar in magnitude to uncertainty of EL.

Values of $\Delta^{(.99)}$ calculated for portfolios (a) - (c) are provided in Table 1. Unfortunately, there is no corresponding estimate of $\Delta^{(.99)}$ from the AIR analysis. Recall that the values of $\Delta^{(.99)}$ in Table 1 reflect uncertainty in the CAT simulation model due to the finite number of historical storms. While a 99% 'confidence level' for EL is reported within the DCR rating of the USAA/Res Re transaction,[Buckley and Orr, 1998] this confidence level reflects process risk only.[Buckley, 1998] This can be understood in the following way: Assuming that the CAT model simulates 10,000 years, then a 1.35% attachment point (e.g., FL) implies that the contract attaches in 135 of the 10,000 years. The Monte Carlo process of having years fall over the attachment point may be modeled as a binomial distribution. Therefore N=10,000, p=1.35%, Np =135 and the coefficient of variation CoV = $\sigma/\mu = (Npq)^{1/2}/(NP) = 8.5\%$. On the other hand, an equivalent CoV can be calculated directly from the quoted values of EL reported for the USAA/Res Re transaction. The value of the 99% confidence level for EL reported by DCR is at 18 basis points above the base case. So CoV = $\sigma/\mu = (0.18/2.33)/0.72 = 11\%$, roughly consistent with the CoV calculated using the binomial assumption. The factor of 2.33 corresponds to the 99th percentile of the unit normal distribution. Values of $\Delta^{(.99)}$ calculated for portfolios (a) through (c) range from 1.9 to 2.3. As comparison, the value of $\Delta^{(.99)}$ calculated using the $EL^{(.99)}$ process risk only value is $\Delta^{(.99)} = 0.25$, implying that confidence intervals reflecting uncertainty due to the finite number of historically observed hurricanes is almost an order of magnitude larger than the magnitude of process risk reported for this transaction.

Both climatological uncertainty and estimation process risk represent sources of imprecision in evaluating EL, or equivalently, CAT bond default rates. A key distinction, however, is that the process risk component arises from and can easily be modified by human intervention, such as changing the number of simulation cycles. Since the standard error of a binomial estimate follows a "square-root of N" law, more (or fewer) simulation cycles will produce greater (or lesser) accuracy. Specifically, if 10,000 simulated years imply a process risk only $\Delta^{(.99)}$ of 0.25, then 40,000 simulated years will imply $\Delta^{(.99)} = 0.125$ and 2,500 years will imply $\Delta^{(.99)} = 0.5$. On the other hand, achieving a halving of the climatological uncertainty $\Delta^{(.99)}$ would require a quadrupling of the historical data – a very difficult, if not impossible, task.

These results can also be compared to a study by Moore [1998], who has analyzed uncertainty in the pricing of three CAT related securities. A jacknife technique[Efron, 1993] is applied to the PCS adjusted historic loss ratios for the years 1956 – 1994. The three analyzed securities are a pure CAT call option (or excess-of-loss contract), an excess-of-loss layer, and a knock-in CAT call. Five different distributions are used in the jacknife technique to parametrize the losses. The results are summarized in Table 2 expressed as the normalized difference between the plug-in estimate and the 99th percentile, given by $\Delta^{(.99)}$ as defined by Eq. 8. For comparison, values of $\Delta^{(.99)}$ for portfolios (a) through (c) from Table 1 are reproduced in this table. While values of $\Delta^{(.99)}$ calculated by Moore are somewhat smaller than those calculated here, the agreement between Moore's results and this work is good. Note that the values of $\Delta^{(.99)}$ calculated both by Moore and in this work reflect the epistemological uncertainty due to the limited historical record. Moore's analysis is based directly on the historical PCS loss data, while the analysis of the present work is based on the effects of the meteorological record on those losses.

V. Comparison to Corporate Bond Volatility

The discussions presented so far reflect an attempt to evaluate uncertainty in hurricane CAT models and to evaluate its effect on CAT securities pricing. But what is the investor to make of this? How does uncertainty in CAT models compare with uncertainty in more traditional securities? A simple proxy for uncertainty in traditional securities is the volatility of one-year corporate bond default rates. This volatility can be calculated from historical default rates and directly compared to the statistic $\Delta^{(.99)}$ calculated in the previous section. The statistic $\Delta^{(.99)}$ is interpreted here as a measure of uncertainty in CAT bond default rates.

Table 3 contains a summary of historical one-year corporate bond default rates for investment grade, speculative grade and all corporate bonds for the years 1970 – 1997.[Keenan, 1998]. R_{mean} and R_{sdev} correspond to the mean and standard deviation of the default rate, respectively, as determined from a lognormal best fit to the historical data. A truncated lognormal distribution is used to fit the investment grade data since some of the yearly default rates are equal to zero. In analogy to Eq. 8, volatility of one-year corporate bond default rates can be characterized by the statistic $\Delta_R^{(.99)}$, where $\Delta_R^{(.99)} = (R^{(.99)} - R_{mean}) / R_{mean}$. The quantity $R^{(.99)}$ is the 99th percentile of the default rate distribution as determined by the lognormal fit to the data. In the case of the investment grade data, the 99th percentile is calculated from R_{mean} and R_{sdev} assuming a normal distribution. The statistic $\Delta_R^{(.99)}$ is given in Table 3 for the three bond classes. Values of R_{mean}, R_{sdev} and $\Delta_R^{(.99)}$ for all bond classes were also calculated assuming normally distributed data; although not shown here, those results are similar to those in Table 3. Values of $\Delta_R^{(.99)}$ range from 3.6 to 4.6, while the equivalent values of $\Delta^{(.99)}$ from Tables 1 and 2 range from 1.9 to 2.3. This comparison indicates that uncertainty in corporate bonds, as measured by volatility of one-year default rates, is greater than the uncertainty in CAT models resulting from the historically limited meteorological data.

Most CAT bonds (and all rated CAT bonds) come with model-based default probability estimates. In the 1997 to early 1998 period, CAT bond excess yields (rate of return less risk free rate less expected loss rate) typically ranged from 300 to 500 basis points (bp), whereas comparable corporate bonds traded with excess yields around 100 bp. At the time, this discrepancy was attributed to a combination of the novelty of CAT bonds (hence limited marketability) and to the uncertainty surrounding the risk assessment, so that this spread should diminish in time.[Lane, 1998] DCR's use of a 99th percentile (process risk) level for rating the Residential Re bond does not have a counterpart in the rating of corporate bonds. This asymmetry is attributed to the novelty of CAT bonds.[Buckley, October 1998] Indeed, the "zero-beta" argument from CAPM theory suggests that the excess yields of CAT bonds should converge to zero, given sufficient liquidity.[Froot,1995] In the latter half of 1998, CAT bond and corporate bond yields did indeed move closer together. However, this was a result of the turmoil and "flight to quality" in bond markets raising the spreads on corporates, and not on any change in CAT bond excess yields, which remained relatively unaffected. [Sandor, 1998]

VI. Conclusions

A simple technique is introduced for developing confidence intervals on the exceedance probability curves produced by hurricane catastrophe simulation models. Confidence intervals calculated with this technique reflect uncertainty resulting from the limited number of historically observed hurricanes. Because this source of uncertainty underlies all hurricane CAT simulation models, this technique will be applicable to all CAT models. Uncertainty in the mean annual hurricane frequency is estimated directly from the historical data. Uncertainty in the severity is estimated from a parametric bootstrap calculation applied to simulated losses. The magnitude of confidence intervals on the EP curve is dominated by uncertainty in the severity. This can be understood by considering that, compared to the uncertainty on the severity as calculated by the parametric bootstrap method, the mean annual hurricane frequency is relatively well defined by the past 100 years of observations.

This technique is used to calculate confidence intervals on the EP curve for three different portfolios. When normalized to the respective base-case EP curves, confidence intervals calculated for the three different portfolios are quite similar, and appear to be roughly constant for return periods greater than approximately 80 years. Values for the 95th percentile of the normalized EP curves range from approximately 1.9 to 2.8.

The results of the parametric bootstrap calculation are placed in the context of a recent CAT bond transaction, and used to calculate uncertainty of the expected loss. The uncertainty of the expected loss is measured as a normalized difference to the base-case EL by the statistic $\Delta^{(.99)}$, and reflects pricing uncertainty of CAT bonds. The statistic $\Delta^{(.99)}$ is also used a measure of uncertainty in CAT bond default rates. The magnitude of $\Delta^{(.99)}$ calculated for three different portfolios is in good agreement with the results of Moore [1998], in which uncertainty in the pricing of CAT securities is calculated from historical PCS insurance losses. In contrast, the magnitude of $\Delta^{(.99)}$ calculated using the results of the parametric bootstrap technique is almost an order of magnitude larger than an equivalent measure of the process risk component calculated in an actual CAT bond transaction.

In order to develop an understanding of how investors might compare the new CAT bonds with more traditional investment vehicles, the volatility in one-year corporate bond default rates is compared to the magnitude of uncertainty in hurricane CAT simulation models. Uncertainty in CAT simulation models is measured as the volatility of CAT bond default rates, represented by the statistic $\Delta^{(.99)}$, calculated using the results of the parametric bootstrap analysis. The results of this comparison suggest that uncertainty in corporate bonds, as measured by the volatility in one-year corporate bond default rates, is greater than an equivalent measure of volatility in CAT bond default rates. While the "zero-beta" argument from CAPM theory suggests that CAT bond excess yields should converge to zero, (as opposed to corporate bonds which must bear a systematic risk premium)[Froot,1995] recent CAT bond excess yields have been even larger than excess yields for comparably rated corporate bonds. The analysis presented in this paper suggests that, CAPM aside, if investors impose risk premia for the uncertainty of bond returns, then corporate bonds and CAT bonds deserve comparable treatment.

The author would like to gratefully acknowledge John Major for his invaluable guidance in this work; and Daniel BarYaacov and Maya Belubekian for their helpful feedback.

References

Buckley, Keith M., Comments made at the 7th International Conference on Funding Catastrophe Protection, Chicago, IL, October 29, 1998

Buckley, Keith M., Duff and Phelps Credit Rating Co., private communication, 1998

Buckley, Keith M. and Jeffrey J. Orr, "Catastrophe Bonds: DCR's Ratings of Insurance-Linked Securitizations," Duff and Phelps Credit Rating Co., Special Report, May 1998

Efron, Bradley and Robert J. Tibshirani, "An Introduction to the Bootstrap", Chapman Hall, New York, 1993

Friedman, Don G., "Natural Hazard Risk Assessment for an Insurance Program", The Geneva Papers on Risk and Insurance, Vol. 9, No. 30, 1984

Froot, Kenneth A., Brian S. Murphy, Aaron B. Stern, and Stephen E. Usher, “The Emerging Asset Class: Insurance Risk,” Guy Carpenter & Company, Inc., July, 1995

Hogg, Robert. V. and Stuart A. Klugman, "Loss Distributions", John Wiley , New York, 1984

Keenan, Sean C., "Historical Default Rates of Corporate Bond Issuers, 1920 – 1997", Moody's Investors Service, 1998

Lane, Morton., “Price, Risk and Ratings for Insurance-Linked Notes,” Sedgwick Lane Financial, May 5, 1998.

Major, John. A., "Uncertainty and Bias: A Frequentist Perspective", Guy Carpenter & Company, Inc., 1998

Moore, James F., "Tail Estimation and Catastrophe Security Pricing: Can We Tell What Target We Hit If We Are Shooting in the Dark?", draft, 1998

Sandor, Richard. L., Comments made at the 7th International Conference on Funding Catastrophe Protection, Chicago, IL, October 28, 1998.

Figure Captions

Figure 1: Results of a bootstrap calculation for the mean annual hurricane frequency with n = 500. The histogram represents the results of the binned bootstrap calculation, and the solid line is a t-distribution corresponding to the empirical mean and variance as defined in the text.

Figure 2: The empirical severity distributions $F_{SE}(\, l)$ determined by a CAT simulation model for three different portfolios (thick solid lines) corresponding to (a) personal, (b) commercial and (c) specialty lines exposure, respectively. Also shown are the best-fits to the empirical severity, $F_{SP}(\, l)$ (thick dashed lines), as well as the parametric bootstrap replication severity distributions $F_{SP}{}^{*b}(\, l)$ plotted for every 10^{th} bootstrap replication (dotted lines).

Figure 3: Plug-in EP curves for the three portfolios (a), (b) and (c) represented as a log-log plot of loss (in dollars) versus return period (in years) (middle solid lines), where return period T = 1/EP. Also shown are percentile EP curves based on the uncertainty in the mean annual frequency and in the severity resulting from the limited historical dataset. The upper solid lines correspond to the 95^{th} percentile curve $EP^{(.95)}$ (*l*) and the lower solid lines corresponds to the 5^{th} percentile curve $EP^{(.05)}$ (*l*). The dashed lines in Fig. 3 correspond to the case where only uncertainty in the mean annual frequency is used to determine a 90% confidence intervals on the EP curve. Calculating confidence intervals on the EP curve based only on uncertainty in the severity produces curves virtually identical to the curves corresponding to the inclusion of both frequency and severity uncertainty

Figure 4: The ratio $L^{(p)}/\hat{L}$ versus return period, where $L^{(p)}$ is the loss percentile and $\hat{L}$ is the plug-in estimate for the loss at the respective return period for percentiles p = 95% (>1) and p = 5% (<1) for the three portfolios (a) through (c).

Tables

	$FL_{\$}$ (\$b)	$DL_{\$}$ (\$b)	$\hat{EL}$(%)	$\Delta^{(.99)}$
	1.35%	*0.55%*		
USAA/Res Re	**1**	**1.5**	**0.72**	**–**
Portfolio (a)	**0.54**	**1.40**	**0.88**	**1.91**
Portfolio (b)	**0.46**	**0.85**	**0.96**	**1.94**
Portfolio (c)	**0.08**	**0.21**	**0.83**	**2.34**

Table 1: Values of the First Loss (FL), Depletion Loss (DL) and Expected Loss (EL) calculated for portfolios (a) through (c) within the context of the USAA/Res Re transaction as described in the text. Details of the USAA/Res Re transaction and modeling results are given for comparison. Also shown is $\Delta^{(.99)}$, a normalized measure of the 99th percentile EL value as determined from a parametric bootstrap analysis.

	Calculation	$\Delta^{(.99)}$
This work:	**Portfolio (a)**	**1.91**
	Portfolio (b)	**1.94**
	Portfolio (c)	**2.34**
Moore:	**Call**	**1.708**
	Spread	**1.252**
	Knockin	**1.827**

Table 2: Comparison of $\Delta^{(.99)}$ calculated in this work for portfolios (a) through (c) along with the results of Moore [1998]. The parameter $\Delta^{(.99)}$ is a measure of the normalized 99th percentile Expected Loss (EL), and is calculated in this work using a parametric bootstrap calculation, and by Moore using a jacknkife technique on historic PCS loss data for three different insurance linked securities.

	Corporate Bond Default Rates (%), 1970 – 1997		
	Investment grade	**Speculative grade**	**All corporates**
R_{mean}	**0.05**	**3.41**	**1.07**
R_{sdev}	**0.10**	**3.15**	**1.24**
$\Delta_R^{(.99)}$	**4.43**	**3.57**	**4.62**

Table 3: Volatility of one-year corporate bond default rates for investment grade, speculative grade and all corporate bonds.[Keenan, 1998] The mean and standard deviation of each bond class is given by R_{mean} and R_{sdev}, respectively. Default rate volatility is quantified by the statistic $\Delta_R^{(.99)}$ as defined in the text.

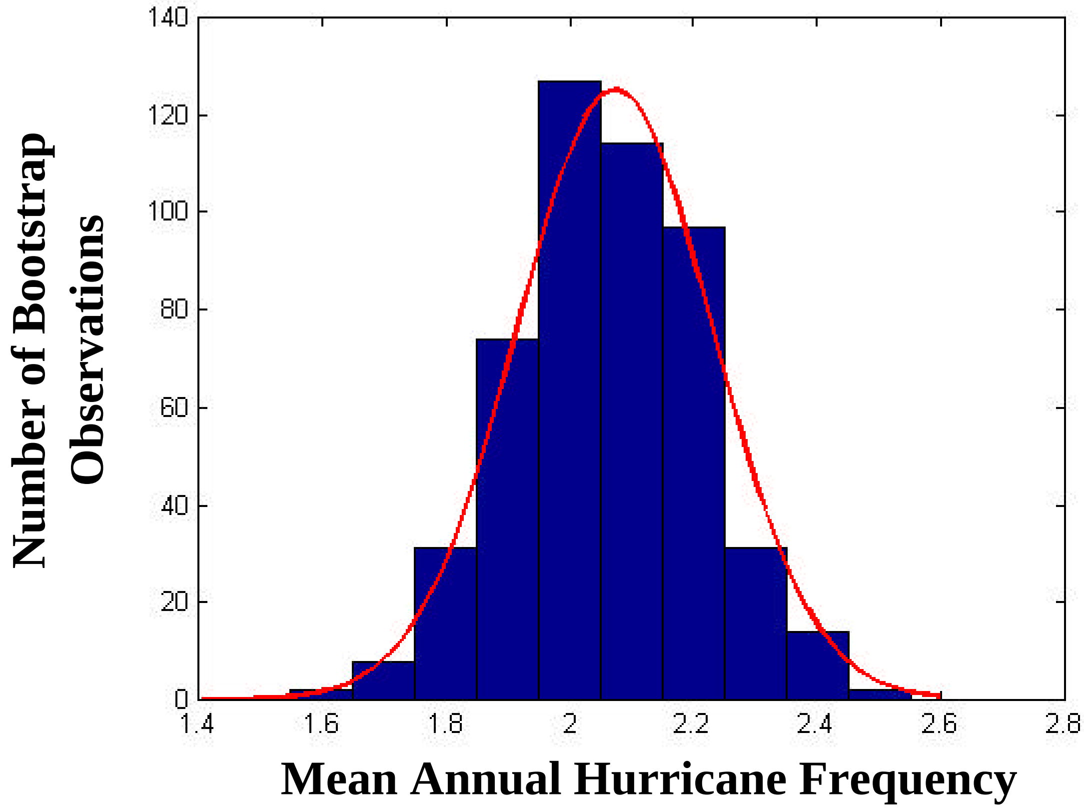

Figure 1: Results of a bootstrap calculation for the mean annual hurricane frequency with n = 500. The histogram represents the results of the binned bootstrap calculation, and the solid line is a t-distribution corresponding to the empirical mean and variance as defined in the text.

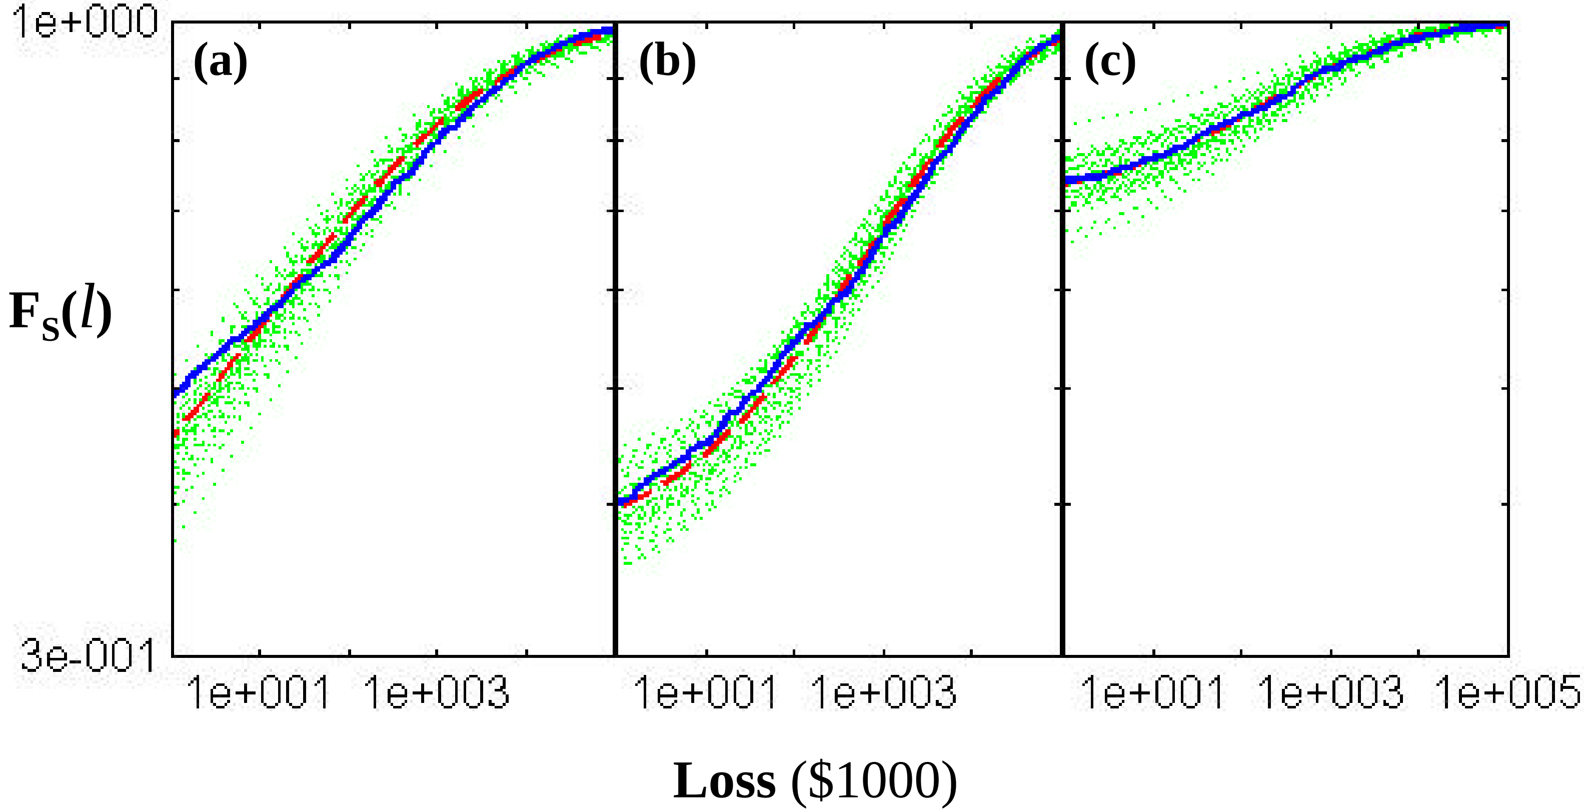

Figure 2: The empirical severity distributions $F_{SE}(l)$ determined by a CAT simulation model for three different portfolios (thick solid lines) corresponding to (a) personal, (b) commercial and (c) specialty lines exposure, respectively. Also shown are the best-fits to the empirical severity, $F_{SP}(l)$ (thick dashed lines), as well as the parametric bootstrap replication severity distributions $F_{SP}{*}^{b}(l)$ plotted for every 10th bootstrap replication (dotted lines).

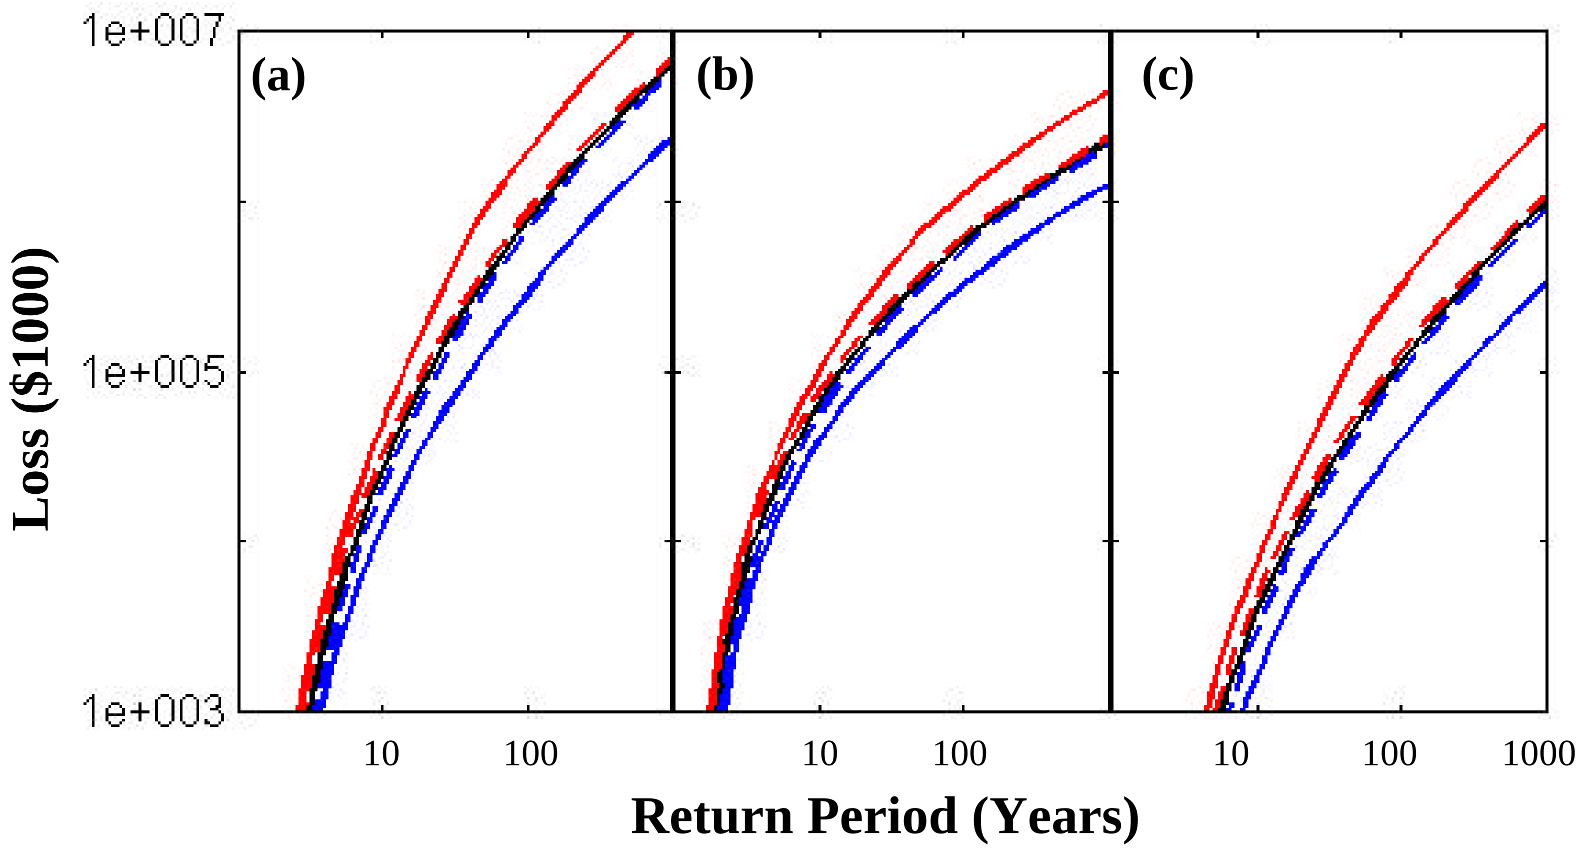

Figure 3: Plug-in EP curves for the three portfolios (a), (b) and (c) represented as a log-log plot of loss (in dollars) versus return period (in years) (middle solid lines), where return period T = 1/EP. Also shown are percentile EP curves based on the uncertainty in the mean annual frequency and in the severity resulting from the limited historical dataset. The upper solid lines correspond to the 95th percentile curve $EP^{(.95)}(l)$ and the lower solid lines corresponds to the 5th percentile curve $EP^{(.05)}(l)$. The dashed lines in Fig. 3 correspond to the case where only uncertainty in the mean annual frequency is used to determine a 90% confidence intervals on the EP curve. Calculating confidence intervals on the EP curve based only on uncertainty in the severity produces curves virtually identical to the curves corresponding to the inclusion of both frequency and severity uncertainty.

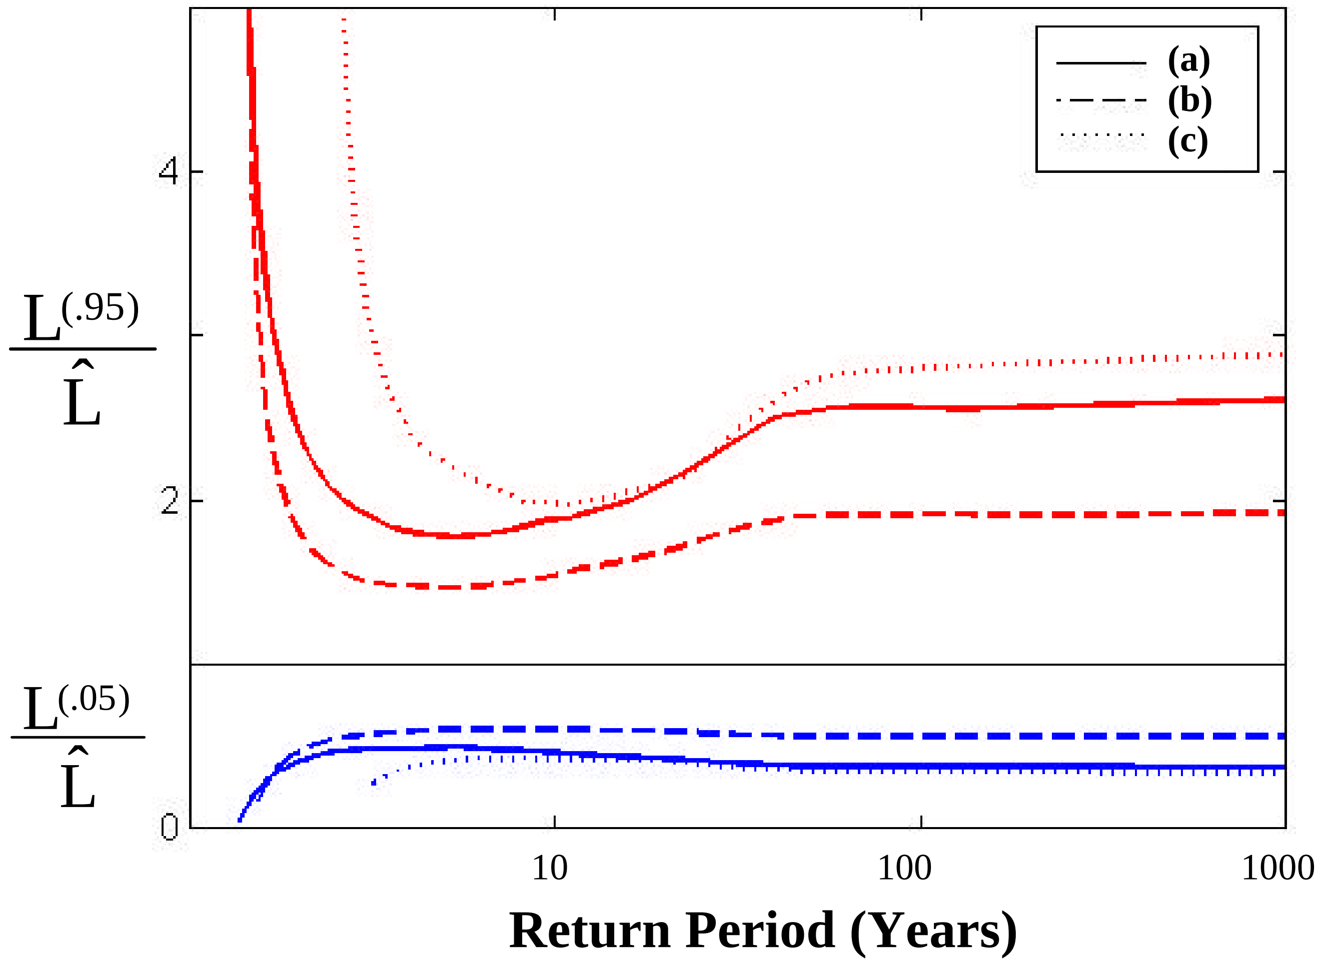

Figure 4: The ratio $L^{(p)}/\hat{L}$ versus return period, where $L^{(p)}$ is the loss percentile and $\hat{L}$ is the plug-in estimate for the loss at the respective return period for percentiles p = 95% (>1) and p = 5% (<1) for the three portfolios (a) through (c).