Taking Uncertainty Into Account: Bias Issues Arising from Parameter Uncertainty in Risk Models

John A. Major, ASA, MAAA

Taking Uncertainty Into Account: Bias Issues Arising from Parameter Uncertainty in Risk Models

by John A. Major, ASA

Given a random variable of interest, a historical sample of its realized values, and the desire to model its possible future values, actuarial training provides many methods for selecting a family of probability models (distributions) and determining specific parameter values that best represent it. But how should one take parameter uncertainty (parameter risk) into account? In particular, uncertainty can lead to bias in estimators commonly used by actuaries. This paper examines the problem of adjusting estimated distributions (risk curves) to remove the undesirable bias effects of parameter risk, and shows several solutions. It goes on, however, to critique the very notion of uncertainty-adjusted risk curves, emphasizing that this is an ambiguous concept. The form of the adjustment depends crucially on details of the specific question being addressed, so much so that an estimator can seem to be simultaneously overestimating and underestimating risk. Parameter uncertainty therefore cannot be "taken into account" in an unequivocal manner. It is recommended that parameter risk be held apart from process risk and presented in terms of confidence intervals; only with that as background – and with great care – should bias corrections be attempted.

0. INTRODUCTION

0.1 Parameter Estimation for DFA

For DFA in particular, the problem of parameter estimation occurs in the process of determining the appropriate method for generating random variables in the simulation of a financial security system. For example, if it is desired to investigate the relative efficacy of various reinsurance alternatives, a simulation can be created that tests the alternatives in a series of hypothetical "future histories" of loss experience. To simulate many realizations of possible future losses – many more than have been observed in the past – it is necessary to first create a model of the probability distribution of losses. Such a model would be based, at the very least, on the loss experience observed in the past.

If one can determine an appropriate cumulative probability distribution function (risk curve) $F_X(X,\theta)$ to associate with the random variable of interest X, then random instances of X can be created by the inverse lookup method: $X = F_X^{-1}(U,\theta)$ where U is a uniformly distributed random variable between 0 and 1. For specific distributions, more efficient techniques are available, but inverse lookup will always work when F can be inverted, either analytically or numerically.

Generated variables X, Y, Z, etc., can be combined in pro forma financial statements or other actuarial calculations to simulate financial results $R=\varphi(X,Y,Z,\ldots)$. After a sufficient number of simulation cycles, the empirical distribution of R values can be used to assess the risk to the financial system, answering questions such as "What value of R is not likely to be exceeded with probability q?" and "What is the probability that R will be greater than (a fixed value) L?"

0.2 Randomness and Uncertainty

"The uncertainty associated with a stochastic model has two distinct sources: the inherent variability of the phenomenon, [and] incomplete knowledge... of the probabilities.... Sometimes these sources of uncertainty are referred to as 'process risk' and 'parameter risk,' respectively. The terms 'risk' and 'uncertainty,' respectively, have also been used...." [Committee on Principles, 1997] In this paper, the terms "randomness" and "uncertainty" are used.

Standard statistical theory, as taught to actuaries, offers many methods for fitting risk models (distributions) to data. With parametric models, there are a variety of techniques for estimating the parameters and assessing the uncertainty in those estimates. What is relatively lacking, however, is advice on how to incorporate uncertainty information into the risk model itself, or more generally, into the advice being given to the user of the risk model.

The predictive approach to probability modeling is one such method for embedding uncertainty (parameter risk) into the (process risk) model for a random variable. The random variable's assumed family of distributions and its parameters are augmented to include variation in the estimation process itself. A familiar example of this is the construction of a prediction interval for a yet-to-be observed time series or regression value. The formula for the variance of the predicted value includes terms for both the residual error (noise) variance and the variance of the estimator for the mean value. Another way of saying this is that the estimated risk curve for the random variable is modified somehow to account for the phenomenon of parameter uncertainty.

It is the purpose of this paper to critique the predictive approach (or indeed any model-embedded approach) to "taking uncertainty into account" in parameter estimation and risk curve construction. In so doing, it will emphasize that this is not an unambiguous operation. The desired form of the risk curve adjustment depends crucially on subtle details of the specific question being addressed, so much so that a risk curve can seem to be simultaneously overestimating and underestimating risk.

0.3 Contents

This paper consists of six parts. The remainder of the introduction discusses previous literature in this area. In particular, a seminal work by Kreps [1997] is summarized. Part 1 discusses estimation and bias in the context of probability distribution parameters and percentiles. To help clarify theory, an exponential example and a lognormal example are worked out in some detail. The lognormal example is the same one used by Kreps [1997]. Part 2 presents some motivation for "adjusting for uncertainty." Taking a Frequentist approach, it casts the issue in terms of a particular type of bias and works out the necessary – predictive – adjustment for the two examples. While Frequentist, it draws strong parallels to the Bayesian approach in Kreps [1997]. Part 3 extends the bias concerns of part 2 in other directions and reveals the existence of an apparent paradox, making the case against adjustment. Part 4 discusses confidence intervals as an alternative to "adjusting for uncertainty." Confidence intervals for parameters, percentiles, and exceedance probabilities are given for the two examples. Part 5 concludes with advice to the DFA practitioner.

0.4 Previous Research

Previous actuarial literature has addressed "parameter uncertainty," but it is sometimes not clear what the term is intended to encompass.

0.4.1 The View from PCAS 1983

Venter [1983] refers to the possibility of modeling "parameter risk" in the context of transformed gamma and beta models for losses where "because of uncertain trend (or other factors) there is substantial uncertainty about the scale parameter λ...." He goes on to suggest putting a gamma distribution on λ^{α} and mixing the loss distribution over λ, as a "practical technique for quantifying this uncertainty." The parameters for the distribution of λ itself can be estimated through percentile matching or, alternatively, an examination of industry or sub-sector loss ratios.

Meyers & Schenker [1983] and Heckman & Meyers [1983] discuss parameter uncertainty in the collective risk model. "Parameter uncertainty can arise from sampling variability and changes... over time.... [or] when some members of the group have different [expectations]." Their model uses a "contagion parameter" c in the claim count distribution and a "mixing parameter" b in the claim severity distribution. Specifically, λ, the expected number of claims (say, from a Poisson distribution), is multiplied by χ, a gamma-distributed random variable with mean 1 and variance c. Z, the claim amount, is divided by β, a gamma-distributed random variable whose inverse has mean 1 and variance b.

Meyers & Schenker [1983] provide three examples of fitting the parameters b and c to empirical data. In the most general form, their model treats r_i years of experience of

insureds $i = 1,\ldots,T$ as manifesting T independent draws of the χ and β random variables. Their equations then estimate b and c through variance components (random effects ANOVA).

Thus, we seem to have three sources of parameter uncertainty which perhaps should be carefully distinguished: sampling error, nonstationarity, and heterogeneity. The recommended mathematical treatment is to interpret uncertainty as a hierarchical random effect. While this method admirably represents nonstationarity and heterogeneity, it does not appear to address sampling error. Sampling error is distinct from heterogeneity; it determines the accuracy with which b, c, λ, etc., can be estimated. The standard errors of the estimates will diminish with increasing numbers of insureds T. The values of b and c themselves, however, will not converge to zero with increasing T.

0.4.2 Kreps 1997

Kreps [1997] discusses parameter uncertainty in normal and lognormal distributions. In his introduction, he states "One of the most ubiquitous sources of parameter uncertainty is the fact that samples in real life are never infinite." Here, he is explicitly addressing sampling error, and develops a theory of predictive distributions "with" parameter uncertainty.[1] He concludes that "the effect of parameter uncertainty is to push probability away from the mean out into the tail." As will be seen below, the

[1] Mathematically, his technique is again to treat uncertainty as a hierarchical random effect, however, with the imprimatur of explicitly Bayesian justifications.

predictive approach can be interpreted as creating percentile estimators that are unbiased in a probabilistic sense.

For a case study, he analyzes Best's reserving data. IBNR is assumed to be distributed lognormally. Based on n=5 years, the maximum likelihood estimates of the mean and standard deviation of X = ln(IBNR) are 23.01923 and 0.06653, respectively. This "point estimate" implies a probability of IBNR exceeding $11.5 billion equal to 1.39%. For Kreps, taking parameter uncertainty into account, "the exact result... is 12.78%. To get to the true 1.39% level, it is necessary to reserve $14.1 billion!"

Subsequent sections will follow through on this example and parallels to Kreps's work will be sketched in more detail.

1. ESTIMATION

This section discusses the estimation of parameters and percentage points. While the estimation of parameters is the usual goal, the theory of point estimation applies equally well to the estimation of functions of the parameters. Because of the typical DFA interest in tail behavior of variables, the estimation of percentiles (specific points on the risk curve) is arguably more important than the estimation of parameters per se. At the very least, the choice of parameter estimation technique should be informed by the effect it has on percentile estimates. *Bias* is defined and illustrated in both parameter and percentile contexts. The concept of a *risk curve* is formally defined and examples are presented. The specific notions of X-unbiased risk curves and estimation techniques are defined and illustrated.

1.1 Estimation of Parameters

While various techniques are available for estimating parameters, we focus here on Maximum Likelihood due to its general applicability and widespread use. Consider a family of probability density functions $g(x;\theta)$ where x is a real variable and θ is a (possibly vector) parameter. Given a sample $\{x_1, x_2, \ldots, x_n\}$, the Maximum Likelihood Estimate (MLE) of the parameter θ is the value $\hat{\theta}$ that maximizes the joint likelihood

$$L(\theta) = \prod_{i=1}^{n} g(x_i;\theta). \tag{1.1}$$

The sampling distribution of $\hat{\theta}$ has (asymptotically, i.e. with large samples) a dispersion matrix equal to the inverse of the matrix of second derivatives (with respect to θ) of the natural log of the likelihood. Thus, standard errors of the MLE may be computed nearly as easily as the estimator itself. In many commonly-used families of distributions, the MLEs are the obvious moment estimators.

For the typical distributions in use by actuaries, MLEs are *asymptotically efficient*. This means that for large samples, they uniformly provide the most accuracy, regardless of the true parameter value. However, they tend not to have strong small-sample justifications [Lehmann, 1983].

1.1.1 The Exponential Case

Consider a random variable X distributed as exponential with scale parameter λ:

$$\Pr\{X \le x\} = F_X(x;\lambda) = 1 - \exp\left(-x/\lambda\right). \tag{1.2}$$

Given a sample $\{x_1, x_2, \ldots, x_n\}$, the likelihood function is given by

$$L(\lambda) = \prod_{i=1}^{n} \frac{1}{\lambda} \cdot \exp(-x_i/\lambda) = \left(\frac{1}{\lambda}\right)^n \cdot \exp\left(-\frac{1}{\lambda} \cdot \sum_{i=1}^{n} x_i\right). \tag{1.3}$$

Differentiating by λ and setting to zero, we can see that the value of λ that maximizes the likelihood is given by

$$T = \hat{\lambda} = \frac{1}{n} \sum_i x_i . \tag{1.4}$$

This is also the same estimator obtained by equating first moments of the theoretical distribution and the sample.

1.1.2 The Normal Case

Consider a random variable X distributed according to the normal cumulative distribution function:

$$\Pr\{X \le x\} = F\left(x, \begin{bmatrix} \mu \\ \sigma^2 \end{bmatrix}\right) = \int_{-\infty}^{x} \frac{1}{\sqrt{2 \cdot \pi}} \cdot \exp\left\{-\frac{1}{2} \cdot \frac{(z-\mu)^2}{\sigma^2}\right\} dz \tag{1.5}$$

The likelihood can be written

$$L\left(\begin{bmatrix} \mu \\ \sigma^2 \end{bmatrix}\right) = \exp\left\{-\frac{n}{2} \cdot \ln(2 \cdot \pi) - n \cdot \ln(\sigma) - \frac{1}{2} \cdot \sum_{i=1}^{n} \frac{(x_i - \mu)^2}{\sigma^2}\right\} \tag{1.6}$$

Differentiating the expression inside the exponential and setting to zero, we get the so-called *likelihood equations*:

$$\frac{\partial \ln(L)}{\partial \mu} = \frac{1}{\sigma^2} \cdot \left(\sum_{i=1}^{n} x_i - n \cdot \mu \right) = 0$$
$$\frac{\partial \ln(L)}{\partial \sigma} = -\frac{n}{\sigma} + \frac{1}{\sigma^3} \cdot \sum_{i=1}^{n} (x_i - \mu)^2 = 0 \qquad (1.7)$$

The solutions, the maximum likelihood estimators, consist of the sample mean and variance, respectively:

$$\hat{\mu} = \frac{1}{n} \sum_i x_i$$
$$\hat{\sigma}^2 = \frac{1}{n} \sum_i (x_i - \hat{\mu})^2 \qquad (1.8)$$

Again, this gets the same result as moment matching. For a lognormal variable Y = exp(X), the sample mean and variance of ln(Y) make up the MLE. This follows from an invariance property of MLEs.

In the example set out in Kreps [1997], we have the log of IBNR modeled as a normal distribution with $\hat{\mu}$ = 23.01923 and $\hat{\sigma}$ = 0.06653 based on n = 5 sample points.

1.2 Estimating a Percentile

Typically, actuarial risk calculations concern themselves with one tail of a distribution. In DFA, the "interesting" or "risky" behavior of the system will often be driven by the upper or lower extreme values of one or more key variables. For example, in the context of reserving, it is common to ask, what level of the loss variable will only be exceeded with specified low probability? This sort of quantity is also known in other financial disciplines as the *value at risk*.

The $100(1-q)^{th}$ percentile X_q of a distribution is given by solving $1-q = F_X(X_q,\theta)$. However, this requires knowing the true value of θ. In practice, we only have some estimator $\hat{\theta}$ of θ, therefore we are left with the problem of constructing estimators of X_q.

1.2.1 The Exponential Case

Given the parameter λ, it is readily determined that $X_q = -\lambda \ln(q)$. This suggests an obvious estimator:

$$\hat{X}_q = -T \cdot \ln(q). \tag{1.9}$$

1.2.2 The Normal Case

For normal variables, $X_q = \mu + z_q\sigma$ where z_q is the $100(1-q)^{th}$ percentage point of the standard normal distribution, e.g., $z_{0.05}=1.645$. Again, this suggests an obvious estimator:

$$\hat{X}_q = \hat{\mu} + z_q \cdot \hat{\sigma} \tag{1.10}$$

For the lognormal, we simply transform by $\hat{Y}_q = \exp(\hat{X}_q)$. Kreps's example notes that the probability of exceeding Y = \$11.5 billion is 1.39% (if the estimated parameters are exactly correct). Equivalently, $\hat{X}_{0.0139} = 23.166$ or $\hat{Y}_{0.0139} = 11.5 \cdot 10^9$.

1.3 Bias in Parameter and Percentile Estimators

Since estimators are themselves random variables, it is meaningful to inquire into their sampling behavior (distributional properties). Imagine there are modelers, m = 1,...,M, each drawing an independent sample $\{x_{1m},\ldots,x_{Nm}\}$ from some fixed distribution.

Each modeler assumes (correctly) the form $F(x;\theta)$ of the distribution, but must estimate the parameter θ based solely on his or her own sample. Each modeler will then, presumably, have a different estimate for θ and some will get closer to the actual value of θ than others.

An estimator S for a quantity $f(\theta)$ is said to be unbiased if

$$E_\theta[S - f(\theta)] = 0 \tag{1.11}$$

where the notation $E_\theta[\,]$ denotes mathematical expectation with respect to the distribution characterized by θ. Note that θ, hence $f(\theta)$, is a fixed number and S is a random variable. In the example of the M modelers, unbiasedness means that the average estimate obtained among modelers, as M gets arbitrarily large, will converge to the true value of the parameter. Unbiasedness is only one property that an estimator may possess, and not having it does not necessarily make an estimator inferior to ones that do.[2]

Note that the definition of unbiasedness applies to estimators of any quantity associated with a distribution, parameters as well as percentiles, exceedance probabilities, etc.

---

[2] "Bias" is such a loaded word that statisticians would have been better off with a more technical term like "expectation neutrality." Alas, we are stuck with the baggage of historical usage.

1.3.1 The Exponential Case

The distribution of T, the MLE for the exponential scale parameter λ, can be shown to be a gamma with scale parameter λ/n and shape parameter n,

$$\Pr\{T \le t\} = \int_0^t \frac{\left(\frac{z}{\lambda/n}\right)^{n-1} \exp\left(-\frac{z}{\lambda/n}\right)}{(n-1)!\,\lambda/n}\,dz. \tag{1.12}$$

The mean of T is therefore λ, and the variance is λ^2/n. T is therefore an unbiased estimator for λ. Because T is unbiased for λ, $\hat{X}_q$ is also unbiased for X_q.

1.3.2 The Normal Case

The sample mean of a normal distribution is distributed as a normal with mean μ and variance σ^2/n, therefore it is unbiased for μ. The sample variance is distributed as σ^2/n times a χ^2(n-1) variable; the MLE for σ is therefore biased. We can distinguish several alternatives. If an unbiased estimate of the *variance* (σ^2) is desired, then we want the familiar

$$\hat{\sigma}_s = \sqrt{\frac{n}{n-1}} \cdot \hat{\sigma} \tag{1.13}$$

This gives us a value of 0.07439 in the Kreps example.

Unbiased estimation of the *standard deviation* (σ) is much less familiar to beginning students of statistics. Lehmann [1983] gives a general form for unbiased σ^k estimation[3] which specializes for k = 1 to:

$$\hat{\sigma}_u = \sqrt{\frac{n}{2}} \cdot \frac{\Gamma\left(\frac{n-1}{2}\right)}{\Gamma\left(\frac{n}{2}\right)} \cdot \hat{\sigma} \tag{1.14}$$

This gives us a value of 0.07911 in the Kreps example.

We may generalize our percentile estimator by considering

$$\hat{X}_q = \hat{\mu} + z_q \cdot \hat{\sigma}_{()} \tag{1.15}$$

where we have a choice of estimators $\hat{\sigma}_{()}$ for σ. Recall that the ML estimator of the 1.39% exceedance point ($z_{0.0139}$ = 2.2) is X = 23.166 translating to an IBNR of Y = exp(23.166) = 11.5 billion.

An unbiased estimator for X_q uses $\hat{\sigma}_{()} = \hat{\sigma}_u$ which yields 23.193, translating to an IBNR of 11.82 billion. This is not unbiased for IBNR, however, because an unbiased X does not imply an unbiased $exp(X)$. This author is not aware of an unbiased estimator for Y_q. We can estimate the magnitude of the bias, however, by noting that if the normal parameters were indeed equal to their ML estimators, then, approximately,

---

[3] Johnson, Kotz, and Balakrishnan [1994] discuss the special case of k=1 and present a simpler approximation.

$$\frac{E\left[\exp(\hat{X}_{q,s})\right]}{Y_q} = \exp\left(\frac{\sigma^2}{2}\cdot\left[\frac{1}{n}+z_q{}^2\cdot\left\{\frac{n-1}{2}\cdot\left(\frac{\Gamma\left(\frac{n-1}{2}\right)}{\Gamma\left(\frac{n}{2}\right)}\right)-1\right\}\right]\right) \tag{1.16}$$

where the subscript s indicates we are using the unbiased estimator for σ. This is only an approximation because it assumes that $\hat{\sigma}_s$ is distributed as a normal variable; for $\sigma<1$, however, it is accurate to within 5%. In our example, for values of σ in the neighborhood[4] of the ML value, the ratio of equation 1.16 is within 1.3% of unity, indicating little bias. However, for larger values of σ, the bias can be substantial.

1.4 The Risk Curve and X-Unbiasedness

We can present the results of many percentage point estimators in graphical form. The locus of points $\{<X_q,q>\}$ is known as the risk curve or exceedance probability (EP) curve. We place the exceedance probability q on the vertical axis and the percentile estimate $\hat{X}_q$ on the horizontal axis. Depending on the range of interest, we may want to plot one or both axes logarithmically. An alternative for the vertical axis is to plot the return period, 1/q, in units of time, e.g., years if the variable represents an annually measured quantity.

---

[4] Specifically, for values of the parameter within a two-tailed 90% confidence interval, as defined in section 4.1.

If, for every q, the percentile estimator $\hat{X}_q$ is unbiased, we say that the risk curve is X-unbiased, or unbiased in the X domain. If a parameter estimation technique leads to an X-unbiased risk curve, we will call it an X-unbiased technique.

1.4.1 The Exponential Case

Having developed T, the locus of points $\langle -T \cdot \ln(q), q \rangle$ is the ML risk curve. This risk curve is unbiased in the X domain. The MLE technique for exponentially-distributed data is thus X-unbiased.

1.4.2 The Normal Case

Depending on which $\hat{\sigma}_{11}$ is used, there are corresponding alternatives for the risk curve. Figure 1 shows the MLE-based curve as a thin solid line and the X-unbiased (approximately Y-unbiased) curve as a thick dotted line. For reference, the target <$1.152 billion, 1.39%> probability point is marked with a box. Note that the two versions of the curve differ markedly. Around the reference point, the difference amounts to $300mm on the dollar axis or 1.7% on the probability axis. The MLE technique for normally distributed data is therefore not X-unbiased, but an X-unbiased alternative, based on equation 1.14, is available.

2. THE CASE FOR ADJUSTMENT

Unbiased estimation in the X (log) or Y (dollar) domain may or may not be appropriate for the decisions to be made in a real application of the theory. For example, while the American Academy of Actuaries [1993] says, "Consideration must also be

given to any [statistical] bias in the reserves or premiums," it doesn't specify in what manner this consideration should be given. This section considers a different sort of bias, leading to the notion of P-unbiasedness, and how that can be achieved through the predictive distribution approach.

2.1 Probabilistic Bias and Predictive Bounds

We can ask a slightly different question about estimators for X_q (equivalently, Y_q), based on the property they purport to represent, namely, an exceedance probability of q: What is the expected value of this probability? In particular, we might like estimators that are "probabilistically unbiased" (P-unbiased) in the sense that

$$E\left[1 - F(\hat{X}_q, \theta)\right] = q. \tag{2.1}$$

Such probabilistically unbiased estimators do exist. They are known as *prediction bounds*, because

$$E\left[F(\hat{X}_q, \theta)\right] = \Pr\left\{X \le \hat{X}_q\right\} \tag{2.2}$$

where X is another draw from the population, independent of the sample upon which the estimator is based. Since X_q is the point satisfying $1\text{-}q = F_X(X_q,\theta)$, if $\hat{X}_q$ is unbiased for X_q, it is natural to assume that the probability of $X > \hat{X}_q$ is also equal to q. This is not generally the case: X-unbiasedness does not imply P-unbiasedness. By establishing the true "predictive probability" of an estimator $\hat{X}_q$

$$\wp(q,\theta) = \Pr\left\{X > \hat{X}_q\right\} = E_\theta\left[1 - F_X(\hat{X}_q;\theta)\right] \tag{2.3}$$

we might be able to solve for an adjusted q* satisfying $\wp(q^*, \theta) = q$. Then, $\hat{X}_{q^*}$ may serve as a P-unbiased estimator for X_q. Other routes are available, also. If, for every q, the percentile estimator $\hat{X}_{q^*}$ is P-unbiased, we say that the risk curve is P-unbiased. If a parameter estimation technique leads to an P-unbiased risk curve, we will call it an P-unbiased technique.

2.1.1 The Exponential Case

The predictive probability for an exponential percentile MLE is independent of the parameter:

$$\Pr\left\{X > \hat{X}_q\right\} = \left(\frac{n}{n - \ln(q)}\right)^n. \tag{2.4}$$

For example, with n = 20 and nominal q = 0.01, the true predictive probability is 0.016.

Inverting the relationship, we get the adjusted q* for a "probabilistically unbiased" $\hat{X}_{q^*}$:

$$q^* = \exp\left(n \cdot \left(1 - q^{-1/n}\right)\right). \tag{2.5}$$

For example, with n = 20 and q = 0.01, the computed q* = 0.006. The adjusted (P-unbiased) risk curve is then the plot of $\left\langle -T \cdot n \cdot \left(1 - q^{-1/n}\right), q \right\rangle$.

Can we find a P-unbiased estimator for the exponential parameter? In other words, can we compute T in such a way that the straightforward $\hat{X}_q$ from equation 1.9 is

P-unbiased? Not in general; there is no solution S to the equation $-S \cdot \ln(q) = -T \cdot n \cdot \left(1 - q^{-1/n}\right)$ that holds for all 0<q<1 simultaneously (although as n increases without bound, $S=T$ is an asymptotic solution). This means that there is no parameter estimation technique *within the exponential distribution* that yields a P-unbiased risk curve.

That is because the predictive distribution for an exponential variable is not an exponential distribution, it is a Pareto! This can be seen by solving $X = -T \cdot n \cdot \left(1 - q^{-1/n}\right)$ for q in terms of X:

$$q = \left(1 + \frac{X}{T \cdot n}\right)^{-n}. \tag{2.6}$$

In summary: to create an X-unbiased risk curve from presumed exponential data, first determine the MLE T of the exponential parameter as in equation 1.4. Then substitute T for λ in equation 1.2. This is not P-unbiased, however, because the true exceedance probability at an estimated percentile is affected by parameter estimation uncertainty. For a P-unbiased risk curve, construct the Pareto distribution corresponding to equation 2.6. Drawing simulated values X from the Pareto instead of the exponential will "take uncertainty into account" in the sense that the true exceedance probabilities of the simulated percentage points will be accurate in expectation.[5]

---

[5] Here, "in expectation" means "averaged over all random samples of data from the same exponential population."

2.1.2 The Normal Case

A prediction bound which a single future, independently selected normal variable will not exceed with probability q is given by:

$$\tilde{X}_q = \hat{\mu} + t_{q,n-1} \cdot \hat{\sigma}_x \cdot \sqrt{1 + \frac{1}{n}} \tag{2.7}$$

where t is the 100(1-q)th percentile of a Student t distribution with n-1 degrees of freedom. For our example of n = 5 and q = 1.39%, we get t = 3.379 and the prediction bound is $\tilde{X}$ = 23.295, corresponding to $13.08 billion.

This must mean the estimators in section 1 are probabilistically biased. Indeed, by setting $\tilde{X}_{q^*} = \hat{X}_q$ (equations 2.7 and 1.15, respectively) we may compute the predictive probability q^* corresponding to the nominal q probability for the estimator $\hat{X}_q$. The ML estimator for the q = 1.39% exceedance point, shown previously to be biased in the log domain, has an expected actual exceedance of q^* = 7.34%. Thus, in probabilistic terms, it is drastically biased – downward – underestimating the tail risk.

What about the (log) unbiased estimator based on $\hat{\sigma}_x$? This is a little better, with expected actual exceedance of q^* = 4.98%, but it is still far from being unbiased in the probabilistic sense.

Figure 2 adds $\langle \tilde{X}_q, q \rangle$ to the plot of risk curves as a dash-dot line. This represents a dramatic increase in estimated risk. Values of IBNR exceeding $1.35 billion, essentially inconceivable according to the MLE and dollar-unbiased curves (20,000-year return period or higher), are now seen as a distinct possibility (100-year return period).

Is there a P-unbiased estimator for σ? As with the exponential case, no. The predictive distribution is from the Student t family, not the normal family (although, again, in the limit as n increases without bound, there is convergence). To create an X-unbiased (or log-unbiased) risk curve from presumed normal data, the methods of section 1.3.2 suffice. For a P-unbiased risk curve, however, one must construct the Student t distribution corresponding to equation 2.7.

2.2 Discussion

A specific family of distributions will lead to a specific form for the predictive distribution. However, there is an approximation method which can bypass the analysis. By sampling the parameters (according to an estimate of their distribution) as well as the object random variable (according to the particular parameter values selected in their most recent draw), one can create a random variable drawn from a mixture.[6] This mixture represents a predictive distribution insofar as it incorporates variability in the random variable (process risk) as well as uncertainty in the parameters (parameter risk).

Making this sort of adjustment – analytically or numerically – is often what is meant by "taking uncertainty into account." Notice correcting this new sort of bias is a matter of increasing an understated (on average) risk. For typical actuarial distributions with decreasing density in the upper tail and small enough q, on average, the true exceedance probability $1 - F_X(\hat{X}_q, \theta)$ for the quoted value of an unbiased estimator $\hat{X}_q$

[6] cf. Venter's recommendation discussed in section 0.4.1.

will be higher than the nominal probability q from which the estimate is developed. The adjusted value $\hat{X}_{q'}$ will therefore be higher (farther up in the tail); this is why it is often claimed that "uncertainty fattens the tails."[7]

Why does this happen? The function F_x is nonlinear in its X argument. Values of an X-unbiased $\hat{X}_q$ deviate from the true value in a balanced fashion between high and low; the average is the true value X_q. However, a deviation on the high side contributes less to the expectation of $F_X(\hat{X}_q, \theta)$ than an equally large deviation on the low side diminishes it, due to the curvature of F_x. Therefore the expectation is not the same as the function evaluated at the true value X_q.

As mentioned in section 0.4.2, Kreps [1997] addresses this issue from a Bayesian perspective. His result for the "true" 1.39% exceedance point is $14.1 billion, about a billion higher than was calculated in section 2.1.2. It is interesting to note that Kreps [1997] summarizes his computations of percentage points with analogous expressions $\hat{\mu} + z \cdot \hat{\sigma}$ involving the MLEs of the parameters. For the MLE of the percentage point, z is the corresponding percentile of a unit normal. For the predictive distribution, Kreps's z is z_{eff}, the percentile from a normal with variance $(n+1)/(n+\tau-4)$, where τ is a parameter defining the "uninformative" Bayesian prior distribution on σ, typically 0 or 1 (he used zero). Since a t distribution with ν degrees of freedom has variance $\nu/(\nu-2)$, the

[7] cf. Kreps's comment, discussed in section 0.4.2.

equivalent Frequentist coefficient $t_{q,n-1} \cdot \sqrt{\frac{n+1}{n-1}}$ (derived from equation 2.7) can be considered analogous to the Bayesian z_{eff} with $\tau = 1$.

Bayesians feel free to treat uncertainty in the parameters on an equal footing with the stochastic behavior of the random variable. Above, we saw how Frequentist mathematics can, in effect, yield the same results. If probabilistically unbiased estimation (or simulation) is the goal, it is appropriate to utilize the predictive distribution, rather than the ML-estimated distribution, to look up percentiles (or generate random variables). This is the Frequentist rationale for "adjusting the risk curve for uncertainty."

3. THE CASE AGAINST ADJUSTMENT

In this section, the search for hidden forms of bias continues. The concept of Q-unbiasedness will be defined. It will be seen that the adjustments of section 2 can lead to worsening of estimator behavior with respect to Q-unbiasedness. Moreover, it will be seen how it is typically impossible to make an adjustment which simultaneously improves the two competing measures of bias.

3.1 Estimating Exceedance Probabilities and Q-Unbiasedness

Rather than divulge a dollar limit X_q corresponding to a given exceedance probability q, we may view a risk curve as telling us a probability Q_L of exceeding a specific threshold L. This might be the perspective, say, in a ruin-theoretic analysis. The decisionmaker could have in mind that $1.152 billion is the most that could be lost

without dire consequences, and might request an estimate of the probability of suffering them. As far as the geometry of the risk curve is concerned, this new situation is simply a matter of entering the graph from a different axis, treating the locus of points as $<L,Q_L>$ rather than $<X_q,q>$.

If an estimator $\hat{Q}_L$ is unbiased, we will say that a risk curve constructed from such estimators is Q-unbiased. If a DFA model aims at constructing risk curves for both X_q and Q_L lookups, then Q-unbiasedness and P-unbiasedness are arguably equally desirable.

A natural point estimator is $\hat{Q}_L = 1 - F_X(L,\hat{\theta})$. Indeed, if $\hat{\theta}$ is the MLE of θ, then $\hat{Q}_L$ is the MLE of Q_L. It should come as no surprise that $E_\theta\left[\hat{Q}_L\right]$ does not in general equal Q_L, again, due to nonlinearity of F_X – this time in its θ argument.

In the two examples it will be seen that, on average, the estimated exceedance probability $\hat{Q}_L$ for the specified loss threshold L will be *higher* than the true probability Q_L. To correct for this bias, an adjusted probability estimate $\hat{Q}_L^*$ will have to be *lower* than the estimate $\hat{Q}_L$ computed from maximum likelihood. Thus, this variety of bias is in the direction of *overstating* the risk, in marked contrast with the case of the previous section, which understated the risk. Q-unbiasedness is not the same as P-unbiasedness. In a sense, they are duals, if not opposites, of each other.

3.1.1 The Exponential Case

The point estimator $\hat{Q}_L$ is obtained from equation 1.2 as exp(-L/T). This is biased, and Johnson, Kotz, and Balakrishnan [1994] give the minimum variance unbiased (MVU) estimator as:

$$\hat{Q}_{LUB} = \left(1 - \frac{L}{n \cdot T}\right)^{n-1}. \tag{3.1}$$

This represents the risk curve as a form of beta distribution. As with P-unbiasedness, there is no estimator of the exponential parameter to make a Q-unbiased exponential risk curve. However (and again, similarly), in the limit as n increases without bound, equation 3.1 approaches an exponential. With $\lambda = 1$, $n = 20$, and $L = 4.605$, the true value of Q_L is 1%; were T to equal λ, this estimator would produce the value 0.69%.

This estimator has the unfortunate property that if L is greater than nT then the estimated exceedance probability is zero, making very-high-tail estimates impractical. By taking a Taylor expansion, we may approximate

$$E_\lambda\left[\hat{Q}_L\right] = E_\lambda\left[1 - F_X(L;T)\right] \approx \exp\left(-\frac{L}{\lambda}\right)\left(1 + \frac{L\left(L - 2\cdot\lambda\right)}{2\cdot n\cdot\lambda^2}\right) \tag{3.2}$$

Unfortunately, the "bias correction" term in this approximation is dependent on the true value of λ, which is unknown. By substituting T for λ, we may compute an approximately unbiased estimate as:

$$\hat{Q}_L{}^* = \exp\left(-\frac{L}{T}\right) \Big/ \left(1 + \frac{L(L-2\cdot T)}{2\cdot n\cdot T^2}\right). \tag{3.3}$$

For tail thresholds L greater than twice the estimated mean T, the denominator is greater than one and the estimated probability is therefore less than the MLE. In this numerical example (n = 20, L = 4.605, and λ = 1), simulation shows this estimator to average 1.1% versus the true 1%. For n = 20, L = 4.605, and T = 1, this adjusted estimator produces 0.77%, about three-fourths of the ML-estimated probability, and 11% higher than the MVU estimator.

3.1.2 The Normal Case

Again we have a variety of estimates

$$\hat{Q}_{()}(x) = 1 - \Phi\left(\frac{x-\hat{\mu}}{\hat{\sigma}_{()}}\right) \tag{3.4}$$

available, depending on the estimator used for σ. Here, Φ is the cumulative (standard) normal probability function corresponding to the integral in equation 1.5. The ML version of this estimator gives us an exceedance probability estimate at L = $11.5 billion ($x$ = ln(L)) of 1.39%.

At this point, readers should not be surprised to learn that the MLE is biased. Lehmann [1983] and Johnson, Kotz, and Balakrishnan [1994] provide the minimum variance unbiased estimator for the exceedance probability of a normal distribution:

$$\hat{Q}_{UB}(x) = 1 - \sqrt{\frac{n}{n-1}} \cdot \frac{\Gamma\left(\frac{n-1}{2}\right)}{\Gamma\left(\frac{1}{2}\right) \cdot \Gamma\left(\frac{n-2}{2}\right)} \cdot \int_S^U \left(1 - \frac{n}{n-1} \cdot z^2\right)^{\frac{n}{2}-2} dz \qquad (3.5)$$

$$where \quad S = \sqrt{\frac{n-1}{n}} \quad and \quad U = \min\left(S, \frac{x - \hat{\mu}}{\sqrt{n} \cdot \hat{\sigma}}\right)$$

Unfortunately, for values of x such that U = S, (which includes our numerical example) this estimator takes on the value of zero. Again, this is likely to be unacceptable in the typical actuarial application.

Alternatively, we can, by numerical integration or simulation, estimate the bias of the MLE (assuming various parameter values). For parameters taking on their MLE values in our example, the expected value for the ML exceedance probability estimator is approximately 1.83%, versus the hypothesized 1.39% – a ratio of 1.3. For other parameter values in the neighborhood[8] of the ML values, this ratio is at least 0.9, usually greater than one, often greater than two, and sometimes greater than 10. This means we should suspect the MLE of being biased high in the situation representing our data, that is, *over*estimating the tail risk. This is in contrast to the MLE percentile estimator, which was biased low, *under*estimating the tail risk.

What about the alternative estimators? Using an unbiased estimator for σ, we get an expected exceedance estimate (again, assuming parameters at the ML values) of 3.17%, high by a factor of 2.3, substantially worse. This is because the unbiased estimate

[8] See footnote 4.

of σ is greater than the ML estimator, decreasing the Z-score, hence the cumulative probability, and hence increasing the exceedance probability.

What about inverting the prediction bound equation? This is the equivalent of "looking up" exceedance probabilities from the predictive distribution. This is worse still, with an expected exceedance estimate of 6.13%, high by a factor of 4.41.

Applying the same strategy as with the exponential distribution, we can take a 2nd order Taylor series approximation to the exceedance probability and express the relative bias as

$$\frac{E_{\mu,\sigma}\left[\hat{Q}_L\right]}{Q_L} = 1 + \frac{1}{Q_L} \cdot \frac{\exp\left(-\frac{1}{2} \cdot z^2\right)}{\sqrt{2 \cdot \pi \cdot n}} \cdot W \tag{3.6}$$

$$\textit{where} \quad W = \left(5 \cdot z - 2 \cdot z^3\right) \cdot \frac{\sqrt{2 \cdot n} \cdot \Gamma\left(\frac{n}{2}\right)}{\Gamma\left(\frac{n-1}{2}\right)} + \left(2 \cdot n - 1\right) \cdot z^3 - 5 \cdot z \cdot n + 3 \cdot z \,.$$

$$\textit{and} \quad z = \frac{L - \mu}{\sigma} \,.$$

As in the case with the exponential, we may substitute the ML estimators for μ and σ, obtaining

$$\hat{Q}_{ML} = 1 - \Phi\left(\frac{x - \hat{\mu}}{\hat{\sigma}}\right) \tag{3.7}$$

$$\hat{Q}^* = \frac{\hat{Q}_{ML}}{1 + \frac{1}{\hat{Q}_{ML}} \cdot \frac{\exp\left(-\frac{1}{2} \cdot \hat{z}^2\right)}{\sqrt{2 \cdot \pi \cdot n}} \cdot \hat{W}}$$

This estimator, while not exactly unbiased, does manage to shrink the bias of the ML estimator, typically by 60-90% in the neighborhood of the ML values.

Figure 3 extends our family of candidate risk curves to include Q^* and Q_{UB}.

3.2 The Paradox

We have seen that to fit a model to data from an assumed distribution, the ML estimators of parameters led us to straightforward construction of risk curves. However, "taking uncertainty into account" in the parameter estimates led us to a profusion of sometimes opposing adjustments.

The ML estimators *under*estimate tail risk in one or two ways. First, the MLE of a normal 100(1-q)th percentile (for small enough q) is, on average, too low. Second, even the unbiased version (or the naturally unbiased estimator in the case of an exponential distribution) still provides "too low" of an estimate because the true exceedance probability of this estimator (the predictive probability) is, on average, greater than the specified amount q.

On the other hand, an MLE of exceedance probability at a (high enough) prespecified threshold is, on average, too high, thereby *over*estimating the tail risk. The substitution of a predictive distribution, corresponding to the probabilistic bias correction for estimating percentiles, is even *more* biased than the MLE.

The search to achieve simultaneous X-, P-, and Q-unbiasedness, even approximately, leads us in conflicting directions.

Consider the implications of this in practice. An actuary has performed a Dynamic Financial Analysis of a client's balance sheet. Numerous sources of random variation in liability and asset values were modeled, each of them having been fit to historical data. After explaining the methodology and walking through various charts

and tables, the actuary summarizes: "There appears to be a 1% chance that your surplus will experience a drop exceeding $1 billion."

In an atypical response, the client might remind the actuary that there is sampling error in the various historical estimates and that actual probability distributions may well be different from the point estimates used in the model. Is this not another source of risk? Should the analysis not be adjusted to "take uncertainty into account?"

They meet a few days later, after the actuary has had a chance to enhance the old "Certainty Model" to include uncertainty adjustments. The following dialog between the client (C) and the actuary (A) ensues.

C: OK, now that you've taken uncertainty into account, what is an unbiased estimate of my 1% exceedance point?

A: It's still $1 billion. That's an unbiased estimate.

C: But isn't it true that exceedance points including sampling error should be higher than exceedance points without?

A: Yes, that makes sense. See, the probability of your experiencing a loss greater than the point the Certainty Model picks out as the 1% point, that is to say, the predictive probability, is actually greater than 1%, so the $1 billion figure is probably too low. A better answer is more like $1.2 billion.

C: So an unbiased estimate of the probability of exceeding $1 billion is actually greater than 1%?

A: No, exceedance estimation in the Certainty Model is biased upwards. An unbiased estimate of the probability of exceeding $1 billion is more like 0.8%.

C: *First you tell me $1 billion is an unbiased estimate for the 1% point. Then you tell me the risk is worse, that the probability is actually greater than 1%. Then you tell me the risk is better, that an unbiased estimate is less than 1%. Now tell me why I shouldn't report you to the Actuarial Board for Counseling and Discipline!*

What, then is the correct response? How is uncertainty to be taken into account?

4. CONFIDENCE INTERVALS

The classical approach to expressing parameter uncertainty is through summaries of the estimator distributions, either moments or selected percentage points. The latter become *confidence intervals* when couched in terms of the probability that the quoted percentage points bracket the true quantity.

Following Hahn & Meeker [1991], we may define a confidence interval as an interval bracketed by two estimators (functions of the sample data) intended to contain an unknown characteristic of the sampled population. Such a characteristic could be a parameter of the distribution, e.g., the mean or standard deviation of a normal distribution, or a function of those parameters, e.g., a percentile or an exceedance probability. The interval will contain the true value of the characteristic with a specified "confidence," e.g., 99%. This can be interpreted by Frequentists in terms of sampling,

because the interval endpoints are random variables.[9] If independent samples were repeatedly drawn, and the interval computed from the samples, then the interval would contain the true value of the characteristic with the specified frequency, e.g., 99% of the time.

4.1 CI for Parameters

4.1.1 The Exponential Case

A 100(1-α)% confidence interval for the exponential parameter θ is given by

$$\hat{\theta}_{\pm} = \frac{T \cdot n}{\gamma_{\left(\frac{1}{2}, 1-\frac{\alpha}{2}, n\right)}} \tag{4.1}$$

where γ is the 100(α/2)th or 100(1-α/2)th percentile from a gamma distribution with shape parameter n. For our numerical example with T=1, the interval is [0.717, 1.509].

4.1.2 The Normal Case

A 100(1-α)% confidence interval for the mean is given by

$$\hat{\mu}_{\pm} = \hat{\mu} \pm t_{1-\frac{\alpha}{2}, n-1} \cdot \frac{\hat{\sigma}}{\sqrt{n-1}} \tag{4.2}$$

---

[9] Interpretation is even easier for Bayesians, because they are free to treat the parameters themselves as random variables.

where t is the 100(1-α/2)th percentile from a Student's t distribution with n-1 degrees of freedom. A 90% confidence interval for the mean in Kreps's reserving example (see previous section) is therefore [22.948, 23.090].

A 100(1-α)% confidence interval for the standard deviation is given by

$$\hat{\sigma}_{\pm} = \hat{\sigma} \cdot \sqrt{\frac{n-1}{\chi^2_{\left(\frac{1}{2}\pm\frac{1-\alpha}{2}\right),(n-1)}}} \tag{4.3}$$

where χ^2 is the 100(α/2)th or 100(1-α/2)th percentile from a chi-square distribution with n-1 degrees of freedom. A 90% confidence interval for σ in our example is [0.0483, 0.1765].

4.2 CI for Percentiles

4.2.1 The Exponential Case

Since the exponential is defined by only one parameter θ, a confidence interval for a q-exceedance percentile can be obtained directly from the confidence interval for the parameter by substituting endpoints:

$$\hat{X}_{q\pm} = \frac{-T \cdot n}{\gamma_{\left(\frac{1}{2}\pm\frac{1-\alpha}{2},n\right)}} \cdot \ln(q) \tag{4.4}$$

4.2.2 The Normal Case

Since the normal is defined by two parameters that must be estimated, the situation is a bit more complex. A 100(1-α)% confidence interval for X_q is given by

$$\hat{X}_{q\pm} = \hat{\mu} + g'_{\left(\frac{1}{2}\pm\frac{1-\alpha}{2}\right),q,n}\cdot\hat{\sigma}_v \tag{4.5}$$

where tables of g' are available in Hahn & Meeker [1991]. More complete tables, as well as the underlying theory based on the noncentral t, are available in Odeh & Owen [1980]. Johnson, Kotz, and Balakrishnan [1994] also give the distribution of $\hat{X}_q$ in terms of the noncentral t. A 90% confidence interval for the 1.39% exceedance point in the example is [23.113, 23.43]. This translates to an IBNR interval of [10.91 billion, 14.98 billion].

4.3 CI for Exceedance Probabilities

4.3.1 The Exponential Case

Again, since the exponential is defined by only one parameter θ, a confidence interval for exceedance probabilities can be obtained directly from the confidence interval for the parameter by substituting endpoints:

$$\hat{Q}_{L\pm} = \exp\left(-\frac{L\cdot\gamma_{\left(\frac{1}{2}\pm\frac{1-\alpha}{2},n\right)}}{T\cdot n}\right). \tag{4.6}$$

4.3.2 The Normal Case

A 100(1-α)% confidence interval for Q is given by

$$\left[\hat{Q}_{L-},\hat{Q}_{L+}\right] = \left[h_{1-\frac{\alpha}{2},\frac{x-\hat{\mu}}{\hat{\sigma}_v},n},\ h_{1-\frac{\alpha}{2},-\frac{x-\hat{\mu}}{\hat{\sigma}_v},n}\right] \tag{4.7}$$

where values of h are tabulated in Odeh & Owen [1980]. For the reserving example, a 90% confidence interval for exceeding $Y = \$11.5$ billion is [0.000617, 0.28351]. This is a stupefyingly large confidence interval, encompassing a factor of 459 between the two extremes. Figure 4 adds the upper and lower 90% confidence risk curves to the previous risk curves.

5. CONCLUSION

This paper examined the general problem of estimating parameters of probability distributions and the specific problem of estimating the actuarially interesting percentage points and exceedance probabilities as captured in the notion of a "risk curve." The choice of risk curve translates directly into the generation of random variables in DFA if the inverse lookup method is used, or, indirectly, as it affects the selection of distributional parameters for other methods. In particular, the paper showed how parameter uncertainty (parameter risk), stemming from sampling variability, can induce bias in estimators. It presented three varieties of bias that a risk curve could exhibit, depending on what aspect of the curve is considered relevant. It demonstrated that, at least in the common examples of exponential and normal/lognormal distributions, there is no way to correct these biases, even approximately, in a single "uncertainty-adjusted" risk curve. The conclusion, that a risk curve estimation procedure can seem to be simultaneously overestimating or underestimating risk, appeared as something of a paradox.

The resolution of this paradox is to examine our intuitive expectations about uncertainty-induced bias.[10] It is not the case that a single "uncertainty-adjusted" curve can replace the "point-estimated" curve, yielding better estimators all the way around. Uncertainty (parameter risk), it seems, cannot be put on a par with randomness (process risk). The problem is inherent in the nature of parameter uncertainty; like a carpet too big for a room, attempts to "flatten it out" in one spot will only make it "bulge up" somewhere else.

The solution that would-be DFA model builders should consider is to make explicit the distinction between uncertainty and randomness by placing (uncertainty) confidence intervals around the (randomness) estimates. For directly fitted distributions, confidence intervals can be calculated as was done in section 4. For DFA outputs, the situation is not so straightforward. The model can be "stress tested" by substituting extreme, but not implausible (see section 4.1), values of the parameters (equivalently, versions of the risk curve) and observing how the results change. More thoroughly, multiple runs, with parameters selected randomly from estimates of their distributions (again, refer to section 4.1) and fixed within each run, can provide multiple versions of the results. These multiple results can be summarized in terms of percentiles of their empirical distribution, giving, in effect, confidence intervals on the model outputs.

---

[10] Bayesians would say that the resolution is to not be concerned about bias; that bias as a statistical concept is problematical per se. I suspect few actuaries would feel totally comfortable with this advice.

After showing a client stress test or confidence interval results, bias can be addressed according to the particular goals of the problem. Given that bias is typically small compared to confidence intervals, a proper appreciation of confidence intervals would tend to dampen concern over the minutiae of bias adjustments.

REFERENCES

American Academy of Actuaries [1993]. *Report on Reserve and Underwriting Risk Factors*, Property/Casualty Risk-Based Capital Task Force.

Committee on Principles [1997]. *General Principles of Actuarial Science Discussion Draft*, Casualty Actuarial Society and Society of Actuaries.

Hahn, Gerald J. and William Q. Meeker [1991]. *Statistical Intervals: A Guide for Practitioners*, New York: John Wiley & Sons.

Heckman, Philip and Glenn Meyers [1983]. "The Calculation of Aggregate Loss Distributions from Claim Severity and Claim Count Distributions," Proceedings of the Casualty Actuarial Society LXX.

Johnson, Norman L., Samuel Kotz, and N. Balakrishnan [1994]. *Continuous Univariate Distributions, volume 1*, New York: John Wiley & Sons.

Kreps, Rodney E. [1997]. "Parameter Uncertainty in (Log)normal Distributions," presented at the 1997 Annual Meeting of the Casualty Actuarial Society.

Lehmann, E. L. [1983]. *Theory of Point Estimation*, New York: John Wiley & Sons.

Meyers, Glenn and Nathaniel Schenker [1983] . "Parameter Uncertainty in the Collective Risk Model," Proceedings of the Casualty Actuarial Society LXX.

Odeh, Robert E. and D. B. Owen [1980]. *Tables for Normal Tolerance Limits, Sampling Plans, and Screening*, New York: Marcel Dekker.

Venter, Gary [1983]. "Transformed Beta and Gamma Distributions and Aggregate Losses," Proceedings of the Casualty Actuarial Society LXX.

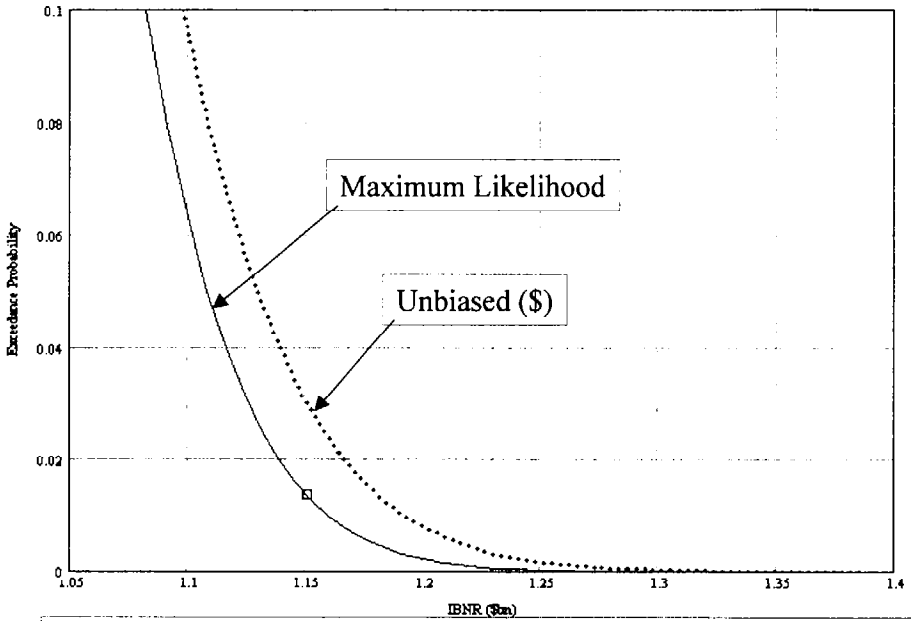

Figure 1: Risk Curves from Lognormal Example (Kreps, 1997)

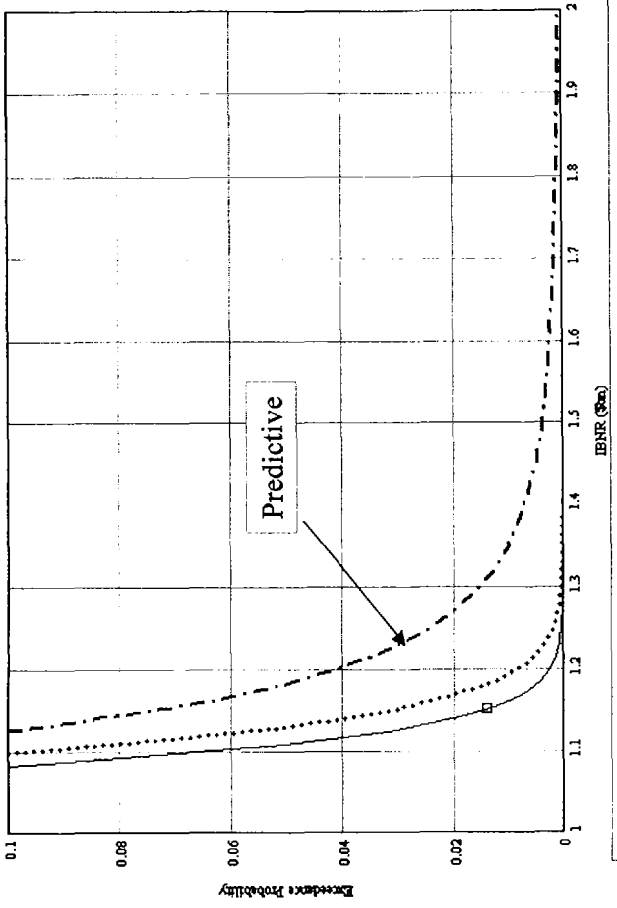

Figure 2: Adding Predictive Distribution Risk Curve

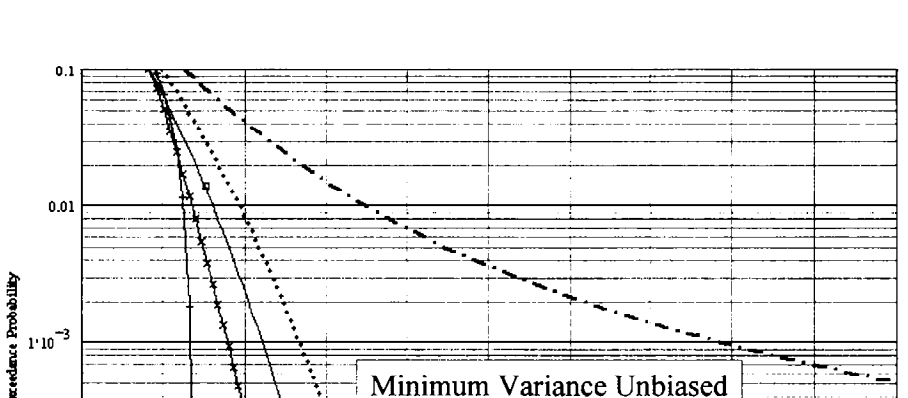

Figure 3: Adding Curves Unbiased for Exceedance

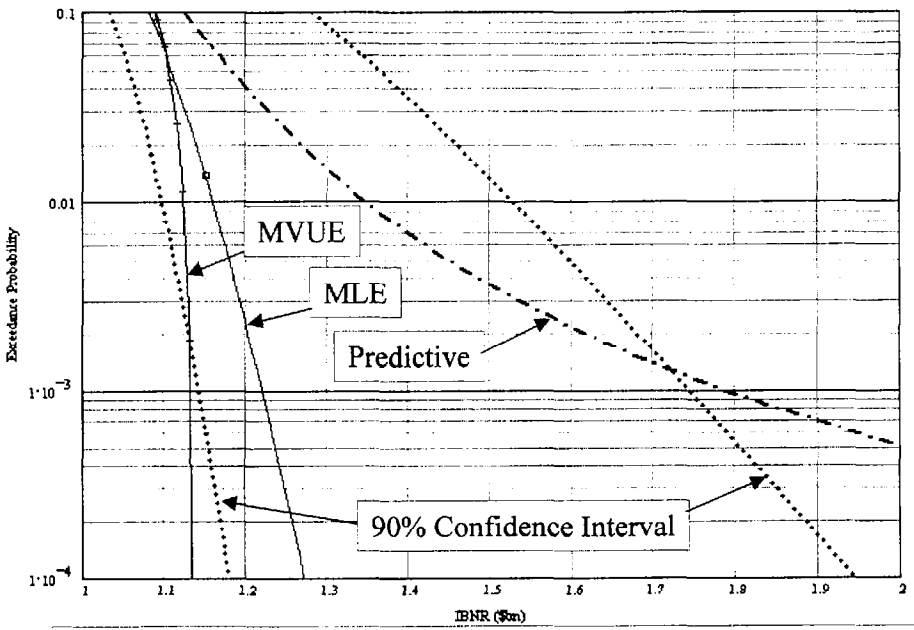

Figure 4: Confidence Intervals Compared to Estimators