

An Introduction to Risk Measures for Actuarial Applications

Mary R Hardy
CIBC Professor of Financial Risk Management
University of Waterloo

1 Introduction

In actuarial applications we often work with loss distributions for insurance products. For example, in P&C insurance, we may develop a compound Poisson model for the losses under a single policy or a whole portfolio of policies. Similarly, in life insurance, we may develop a loss distribution for a portfolio of policies, often by stochastic simulation.

Profit and loss distributions are also important in banking, and most of the risk measures we discuss in this note are also useful in risk management in banking. The convention in banking is to use profit random variables, that is Y where a loss outcome would be $Y < 0$. The convention in insurance is to use loss random variables, $X = -Y$. In this paper we work exclusively with loss distributions. Thus, all the definitions that we present for insurance losses need to be suitably adapted for profit random variables.

Additionally, it is usually appropriate to assume in insurance contexts that the loss X is non-negative, and we have assumed this in Section 2.5 of this note. It is not essential however, and the risk measures that we describe can be applied (perhaps after some adaptation) to random variables with a sample space spanning any part of the real line.

Having established a loss distribution, either parametrically, non-parametrically, analytically or by Monte Carlo simulation, we need to utilize the characteristics of the distribution for pricing, reserving and risk management. The risk measure is an important tool in this

process.

A risk measure is a functional mapping a loss (or profit) distribution to the real numbers. If we represent the distribution by the appropriate random variable \mathbf{X} , and let \mathcal{H} represent the risk measure functional, then

$$\mathcal{H} : X \rightarrow \mathbb{R}$$

The risk measure is assumed in some way to encapsulate the risk associated with a loss distribution.

The first use of risk measures in actuarial science was the development of **premium principles**. These were applied to a loss distribution to determine an appropriate premium to charge for the risk. Some traditional premium principle examples include

The expected value premium principle The risk measure is

$$\mathcal{H}(X) = (1 + \alpha)E[X] \quad \text{for some } \alpha \geq 0$$

The standard deviation premium principle Let $V[X]$ denote the variance of the loss random variable, then the standard deviation principle risk measure is:

$$\mathcal{H}(X) = E[X] + \alpha \sqrt{V[X]} \quad \text{for some } \alpha \geq 0$$

The variance premium principle $\mathcal{H}(X) = E[X] + \alpha V[X]$ for some $\alpha \geq 0$

More premium principles are described in Gerber (1979) and Bühlmann (1970). Clearly, these measures have some things in common; each generates a premium which is bigger than the expected loss. The difference acts as a cushion against adverse experience. The difference between the premium and the mean loss is the premium loading. In the standard deviation and variance principles, the loading is related to the variability of the loss, which seems reasonable.

Recent developments have generated new premium principles, such as the PH-transform (Wang (1995, 1996)), that will be described below. Also, new applications of risk measures have evolved. In addition to premium calculation, we now use risk measures to determine economic capital – that is, how much capital should an insurer hold such that the uncertain

future liabilities are covered with an acceptably high probability? Risk measures are used for such calculations both for internal, risk management purposes, and for regulatory capital, that is the capital requirements set by the insurance supervisors.

In addition, in the past ten years the investment banking industry has become very involved in the development of risk measures for the residual day to day risks associated with their trading business. The current favorite of the banking industry is Value-at-Risk, or VaR, which we will describe in more detail in the next section.

2 Risk Measures for Capital Requirements

2.1 Example Loss Distributions

In this Section we will describe some of the risk measures in current use. We will demonstrate the risk measures using three examples:

- A loss which is normally distributed with mean 33 and standard deviation 109.0
- A loss with a Pareto distribution with mean 33 and standard deviation 109.0
- A loss of $1000 \max(1 - S_{10}, 0)$, where S_{10} is the price at time $T = 10$ of some underlying equity investment, with initial value $S_0 = 1$. We assume the equity investment price process, S_t follows a lognormal process with parameters $\mu = 0.08$ and $\sigma = 0.22$. This means that $S_t \sim \text{lognormal}(\mu t, \sigma^2 t)$. This loss distribution has mean value 33.0, and standard deviation 109.0¹. This risk is a simplified version of the put option embedded in the popular ‘variable annuity’ contracts.

Although these loss distributions have the same first two moments, the risks are actually very different. In Figure 1 we show the probability functions of the three loss distributions in the same diagram; in the second plot we emphasize the tail of the losses. The vertical line indicates the probability mass at zero for the put option distribution.

¹We are not assuming any hedging of the risk. This is the ‘naked’ loss.





Figure 1: Probability density functions for the example loss distributions.

2.2 Value At Risk – the Quantile Risk Measure

The Value at Risk, or VaR risk measure was actually in use by actuaries long before it was reinvented for investment banking. In actuarial contexts it is known as the quantile risk measure or quantile premium principle. VaR is always specified with a given confidence level α – typically $\alpha=95\%$ or 99% .

In broad terms, the α -VaR represents the loss that, with probability α will not be exceeded. Since that may not define a unique value, for example if there is a probability mass around the value, we define the α -VaR more specifically, for $0 \leq \alpha \leq 1$, as

$$\mathcal{H}[L] = Q_\alpha = \min \{Q : \Pr[L \leq Q] \geq \alpha\} \quad (1)$$

For continuous distributions this simplifies to Q_α such that

$$\Pr[L \leq Q_\alpha] = \alpha. \quad (2)$$

That is, $Q_\alpha = F_L^{-1}(\alpha)$ where $F_L(x)$ is the cumulative distribution function of the loss random variable L .

The reason for the ‘min’ term in the definition (1) is that for loss random variables that are discrete or mixed continuous and discrete, we may not have a value that exactly matches equation (2). For example, suppose we have the following discrete loss random variable:

$$L = \begin{cases} 100 & \text{with probability } 0.005 \\ 50 & \text{with probability } 0.045 \\ 10 & \text{with probability } 0.10 \\ 0 & \text{with probability } 0.85 \end{cases} \quad (3)$$

From this we can construct the following table:

x	$\Pr[L \leq x]$
100	1.00
50	0.995
10	0.95
0	0.85

Now, if we are interested in the 99% quantile, there is no value Q for which $\Pr[L < Q] = 0.99$. So we choose the smallest value for the loss that gives at least a 99% probability

that the loss is smaller – that is we choose a VaR of 50. This is the smallest number that gives has the property that the loss will be smaller with at least 99% probability. That is,

$$50 = \min\{Q : \Pr[L < Q] \geq 0.99\}$$

corresponding to definition (1).

Exercise: What are the 95%, 90% and 80% quantile risk measures for this discrete loss distribution?

Solution: 10; 10; 0.

We can easily calculate the 95% and 99% risk measures for the three example loss distributions.

Example 1. Normal($\mu = 33$, $\sigma = 109$) Loss

Since the loss random variable is continuous, we simply seek the 95% and 99% quantiles of the loss distribution – that is, the 95% quantile risk measure is $Q_{0.95}$ where

$$\begin{aligned} \Pr[L \leq Q_{0.95}] &= 0.95 \\ \text{i.e. } \Phi\left(\frac{Q_{0.95} - 33}{109}\right) &= 0.95 \\ \Rightarrow \left(\frac{Q_{0.95} - 33}{109}\right) &= 1.6449 \\ \Rightarrow Q_{0.95} &= \$212.29 \end{aligned}$$

Exercise: Calculate the 99% quantile risk measure for this loss distribution.

Answer: \$286.57

Example 2. Pareto Loss

Using the parameterization of Klugman, Panjer and Willmot (2004), (but changing the notation slightly to avoid confusion with too many α 's) the density and distribution

functions of the the Pareto distribution are

$$f_L(x) = \frac{\gamma \theta^\gamma}{(\theta + x)^{\gamma+1}}$$

$$F_L(x) = 1 - \left(\frac{\theta}{\theta + x} \right)^\gamma$$

Matching moments, given the mean and variance of 33 and 109^2 , we have $\theta = 39.660$ and $\gamma = 2.2018$. The 95% quantile risk measure is $Q_{0.95}$ where

$$\Pr [L \leq Q_{0.95}] = 0.95$$

that is $F_L(Q_{0.95}) = 0.95$

$$\Rightarrow 1 - \left(\frac{\theta}{\theta + Q_{0.95}} \right)^\gamma = 0.95$$

$$\Rightarrow Q_{0.95} = \$114.95$$

Exercise: Calculate the 99% quantile risk measure for this loss distribution.

Answer: \$281.48

Example 3. Lognormal Put Option:

We first find out whether the quantile risk measure falls in the probability mass at zero. The probability that the loss is zero is

$$\Pr[L = 0] = \Pr[S_{10} > 1] = 1 - \Phi \left(\frac{\log(1) - 10\mu}{\sqrt{10}\sigma} \right) = 0.8749 \quad (4)$$

So, both the 95% and 99% quantiles lie in the continuous part of the loss distribution.

The 95% quantile risk measure is $Q_{0.95}$ such that:

$$\Pr[L \leq Q_{0.95}] = 0.95$$

$$\Leftrightarrow \Pr[1000(1 - S_{10}) \leq Q_{0.95}] = 0.95$$

$$\Leftrightarrow \Pr \left[S_{10} > \left(1 - \frac{Q_{0.95}}{1000} \right) \right] = 0.95$$

$$\Leftrightarrow \Phi \left(\frac{\log(1 - \frac{Q_{0.95}}{1000}) - 10\mu}{\sqrt{10}\sigma} \right) = 0.05$$

$$\Leftrightarrow Q_{0.95} = \$291.30$$

Exercise: Calculate the 99% quantile risk measure for this loss distribution.

Answer: \$558.88

We note that the 95% quantile of the loss distribution is found by assuming the underlying stock price falls at the 5% quantile of the stock price distribution, as the put option liability is a decreasing function of the stock price process.

For more complex loss distributions, the quantile risk measure may be estimated by Monte Carlo simulation.

2.3 Conditional Tail Expectation

The quantile risk measure assesses the ‘worst case’ loss, where worst case is defined as the event with a $1 - \alpha$ probability. One problem with the quantile risk measure is that it does not take into consideration what the loss will be if that $1 - \alpha$ worst case event actually occurs. The loss distribution above the quantile does not affect the risk measure. The Conditional Tail Expectation (or CTE) was chosen to address some of the problems with the quantile risk measure. It was proposed more or less simultaneously by several research groups, so it has a number of names, including Tail Value at Risk (or Tail-VaR), Tail Conditional Expectation (or TCE) and Expected Shortfall.

Like the quantile risk measure, the CTE is defined using some confidence level α , $0 \leq \alpha \leq 1$. As with the quantile risk measure, α is typically 90%, 95% or 99%.

In words, the CTE is the expected loss given that the loss falls in the worst $(1 - \alpha)$ part of the loss distribution.

The worst $(1 - \alpha)$ part of the loss distribution is the part above the α -quantile, Q_α . If Q_α falls in a continuous part of the loss distribution (that is, not in a probability mass) then

we can interpret the CTE at confidence level α , given the α -quantile risk measure Q_α , as

$$CTE_\alpha = E[L|L > Q_\alpha] \quad (5)$$

This formula does not work if Q_α falls in a probability mass, that is, if there is some $\epsilon > 0$ such that $Q_{\alpha+\epsilon} = Q_\alpha$. In that case, if we consider only losses strictly greater than Q_α , we are using less than the worst $(1 - \alpha)$ of the distribution; if we consider losses greater than or equal to Q_α , we may be using more than the worst $(1 - \alpha)$ of the distribution. We therefore adapt the formula of Equation (5) as follows

Define $\beta' = \max\{\beta : Q_\alpha = Q_\beta\}$. Then

$$CTE_\alpha = \frac{(\beta' - \alpha)Q_\alpha + (1 - \beta') E[L|L > Q_\alpha]}{1 - \alpha} \quad (6)$$

The formal way to manage the expected value in the tail for a general distribution is to use distortion functions, which we introduce in the next section.

The outcome of equation (6) will be the same as equation (5) when the quantile does not fall in a probability mass. In both cases we are simply taking the mean of the losses in the worst $(1 - \alpha)$ part of the distribution, but because of the probability mass at Q_α , some of that part of the distribution comes from the probability mass.

The CTE has become a very important risk measure in actuarial practice. It is intuitive, easy to understand and to apply with simulation output. As a mean, it is more robust with respect to sampling error than the quantile. The CTE is used for stochastic reserves and solvency for Canadian and US equity-linked life insurance.

It is worth noting that, since the CTE_α is the mean loss given that the loss lies above the VaR at level α , then a choice of, say, a 95% CTE is generally considerably more conservative than the 95% VaR.

In general, if the loss distribution is continuous (at least for values greater than the relevant quantile), with probability function $f(y)$ then Equation (5) may be calculated as:

$$CTE_\alpha = \frac{1}{1 - \alpha} \int_{Q_\alpha}^{\infty} y f(y) dy \quad (7)$$

For a loss $L \geq 0$, this is related to the limited expected value for the loss as follows:

$$\begin{aligned} CTE_\alpha &= \frac{1}{1-\alpha} \int_{Q_\alpha}^{\infty} y f(y) dy \\ &= \frac{1}{1-\alpha} \left\{ \int_0^{\infty} y f(y) dy - \int_0^{Q_\alpha} y f(y) dy \right\} \end{aligned}$$

Now we know from Klugman, Panjer and Willmot (2004) that the limited expected value function is

$$\begin{aligned} E[L \wedge Q_\alpha] &= E[\min(L, Q_\alpha)] = \int_0^{Q_\alpha} y f(y) dy + Q_\alpha(1 - F(Q_\alpha)) \\ &= \int_0^{Q_\alpha} y f(y) dy + Q_\alpha(1 - \alpha) \end{aligned}$$

So we can re-write the CTE for the continuous case as:

$$\begin{aligned} CTE_\alpha &= \frac{1}{1-\alpha} \{E[L] - (E[L \wedge Q_\alpha] - Q_\alpha(1 - \alpha))\} \\ &= Q_\alpha + \frac{1}{1-\alpha} \{E[L] - (E[L \wedge Q_\alpha])\} \end{aligned}$$

Example 1: Discrete Loss Distribution

We can illustrate the ideas here with a simple discrete example. Suppose X is a loss random variable with probability function:

$$X = \begin{cases} 0 & \text{with probability } 0.9 \\ 100 & \text{with probability } 0.06 \\ 1000 & \text{with probability } 0.04 \end{cases} \quad (8)$$

Consider first the 90% CTE. The 90% quantile is $Q_{0.90} = 0$; also, for any $\epsilon > 0$, $Q_{0.90+\epsilon} > Q_{0.90}$, so the CTE is simply

$$CTE_{0.90} = E[X|X > 0] = \frac{(0.06)(100) + (0.04)(1000)}{0.10} = 460$$

that is, 460 is the mean loss given that the loss lies in the upper 10% of the distribution.

Now consider the 95% CTE. Now $Q_{0.96} = Q_{0.95} = 100$, so to get the mean loss in the

upper 5% of the distribution, we use equation (6), with $\beta' = 0.96$, giving

$$CTE_{0.95} = \frac{(0.01)(100) + (0.04)(1000)}{0.05} = 820$$

Example 2: Normal ($\mu = 33$, $\sigma = 109$) example.

The loss is continuous, so the 95% CTE is $E[L|L > Q_{0.95}]$. Let $\phi(z)$ denote the p.d.f. of the standard normal distribution – that is,

$$\phi(z) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2}$$

then

$$\begin{aligned} CTE_\alpha &= E[L|L > Q_\alpha] \\ &= \frac{1}{1-\alpha} \int_{Q_\alpha}^{\infty} \frac{y}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2}\left(\frac{y-\mu}{\sigma}\right)^2} dy \end{aligned}$$

Let $z = \frac{y-\mu}{\sigma}$ then

$$\begin{aligned} CTE_\alpha &= \frac{1}{1-\alpha} \int_{\frac{Q_\alpha-\mu}{\sigma}}^{\infty} \frac{\sigma z + \mu}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2} dz \\ &= \frac{1}{1-\alpha} \left\{ \int_{\frac{Q_\alpha-\mu}{\sigma}}^{\infty} \frac{\sigma z}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2} dz + \mu \int_{\frac{Q_\alpha-\mu}{\sigma}}^{\infty} \phi(z) dz \right\} \end{aligned}$$

Substituting $u = z^2/2$ in the first integral, and noting that the second integral is

$$\mu \left(1 - \Phi \left(\frac{Q_\alpha - \mu}{\sigma} \right) \right) = \mu(1 - \alpha)$$

gives the CTE formula for the Normal distribution as

$$CTE_\alpha = \mu + \frac{\sigma}{1-\alpha} \phi \left(\frac{Q_\alpha - \mu}{\sigma} \right) \tag{9}$$

So, the 95% CTE for the $N(33, 109^2)$ loss distribution is \$257.83, and the 99% CTE is 323.52

Exercise: Derive the formula for the CTE for a Pareto loss distribution, and calculate the 95% and 99% CTE for the example distribution.

Answer:

$$CTE_\alpha = \frac{\theta}{\gamma - 1} + Q_\alpha \frac{\gamma}{\gamma - 1} \quad \text{for}$$

where $\theta > 0$, $\gamma > 1$. In the example, the 95% CTE is \$243.60 and the 99% CTE is \$548.70.

Example 3: The lognormal Put Option example

The 5% worst case for the lognormal put option liability corresponds to the lower 5% of the lognormal distribution of the underlying stock price at maturity, S_{10} . Let Q_α denote the α -quantile of the S_{10} distribution. Then the 95% CTE is

$$\begin{aligned} CTE_{0.95} &= \frac{1}{0.05} \int_0^{Q_{0.05}} 1000(1-y) \frac{1}{\sqrt{2\pi\sigma y}} \exp\left\{-\frac{1}{2}\left(\frac{\log(y)-\mu}{\sigma}\right)^2\right\} dy \\ &= \frac{1000}{0.05} \left\{ \int_0^{Q_{0.05}} \frac{1}{\sqrt{2\pi\sigma y}} \exp\left(-\frac{1}{2}\left(\frac{\log(y)-\mu}{\sigma}\right)^2\right) dy \right. \\ &\quad \left. - \int_0^{Q_{0.05}} \frac{y}{\sqrt{2\pi\sigma y}} \exp\left(-\frac{1}{2}\left(\frac{\log(y)-\mu}{\sigma}\right)^2\right) dy \right\} \end{aligned}$$

Now, the first term in the $\{ \}$ is simply $F_S(Q_{0.05}) = 0.05$. The second term can be simplified by substituting

$$z = \frac{\log(y) - \mu - \sigma^2}{\sigma}$$

so that the integral simplifies to

$$e^{\mu+\sigma^2/2} \int_{-\infty}^{(\log(Q_{0.05})-\mu-\sigma^2)/\sigma} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2} dz = e^{\mu+\sigma^2/2} \Phi\left(\frac{\log(Q_{0.05}) - \mu - \sigma^2}{\sigma}\right)$$

So we have:

$$\begin{aligned} CTE_{0.95} &= \frac{1000}{0.05} \left\{ (0.05) - e^{\mu+\sigma^2/2} \Phi\left(\frac{\log(Q_{0.05}) - \mu - \sigma^2}{\sigma}\right) \right\} \\ &= 1000 \left\{ 1 - \frac{e^{\mu+\sigma^2/2}}{0.05} \Phi\left(\frac{\log(Q_{0.05}) - \mu - \sigma^2}{\sigma}\right) \right\} \end{aligned}$$

We might note that $(\log(Q_{.05}) - \mu)/\sigma = \Phi^{-1}(0.05) = -1.6445$.

Using this formula for the example loss distribution, we find that the 95% CTE is \$454.14 and the 99% CTE is \$644.10.

Exercise: Calculate the 80% CTE for the put option example.

Answer: The 80% quantile lies in the probability mass, so we must use equation (6) for the CTE. Here $\beta' = 0.8749$ (from (4) above); $Q_{0.80} = Q_{0.8749} = 0$. The CTE is

$$CTE_{0.80} = \left(\frac{1 - 0.8749}{1 - 0.8} \right) E[L|L > 0] = \$165.15$$

Exercise: Derive formulae for the α -quantile and α -CTE risk measures for a lognormal loss random variable, with parameters μ and σ .

Answer:

$$Q_\alpha = \exp \left\{ \Phi^{-1}(\alpha) \sigma + \mu \right\}$$

$$CTE_\alpha = \frac{e^{\mu + \sigma^2/2}}{1 - \alpha} \left(1 - \Phi \left(\frac{\log(Q_\alpha) - \mu - \sigma^2}{\sigma} \right) \right)$$

2.4 Some comments on Quantile and CTE risk measures

1. Clearly $CTE_\alpha \geq Q_\alpha$ with equality only if Q_α is the maximum value of the loss random variable.
2. $CTE_{0\%}$ is the mean loss.
3. $Q_{0\%}$ is the minimum loss; $Q_{50\%}$ is the median loss.
4. In Figure 2 we show the quantile and CTE risk measures for all three example loss distributions, for all values of α between zero and one. In the top diagram, the quantile risk measures at $\alpha = 0$ are zero for the Pareto and put option examples, because that is the minimum loss. The Normal example allows profits, so the lower quantile risk measures are negative. The put option quantile risk measure remains at zero for all $\alpha < 0.875$.

In the lower diagram the CTE's all meet at the left side, as all the examples have the same mean loss. At the far right side, the heavy tail of the Pareto distribution

becomes apparent, as the extreme CTE's show very high potential losses. The maximum value of α shown in the plot is 0.999; at this level the Pareto example CTE is \$1634, the Normal example is \$400 and the put option example is \$783.

5. In some cases negative values may be excluded from the calculation, in which case the Normal example quantile risk measure would be zero for $\alpha < 0.38$, and would follow the same path as shown in the upper diagram after that.

For the CTE, the early values for the Normal example would increase, as the CTE would be defined for the loss random variable L as

$$E[\max(L, 0) | L > Q_\alpha].$$

2.5 Distortion Risk Measures

Distortion risk measures are defined using the survival function (decumulative distribution function) for the loss, $S(x) = 1 - F(x)$. We only consider non-negative losses, that is $L \geq 0$ (these methods can be adapted for profit/loss distributions).

The distortion risk measures are those that can be expressed in the form:

$$\mathcal{H}(X) = \int_0^\infty g(S(x)) dx \tag{10}$$

where $g()$ is an increasing function, with

$$g(0) = 0 \quad \text{and} \quad g(1) = 1.$$

The function $g()$ is called the distortion function. The method works by reassigning probabilities such that the worst outcomes are associated with an artificially inflated probability. The function $g(S(x))$ is a risk-adjusted survival function.

The quantile and CTE risk measures both fall into the class of distortion risk measures. They are by far the most commonly used distortion measures for capital adequacy, but others are also seen in practice, particularly for premium setting in property and casualty insurance.

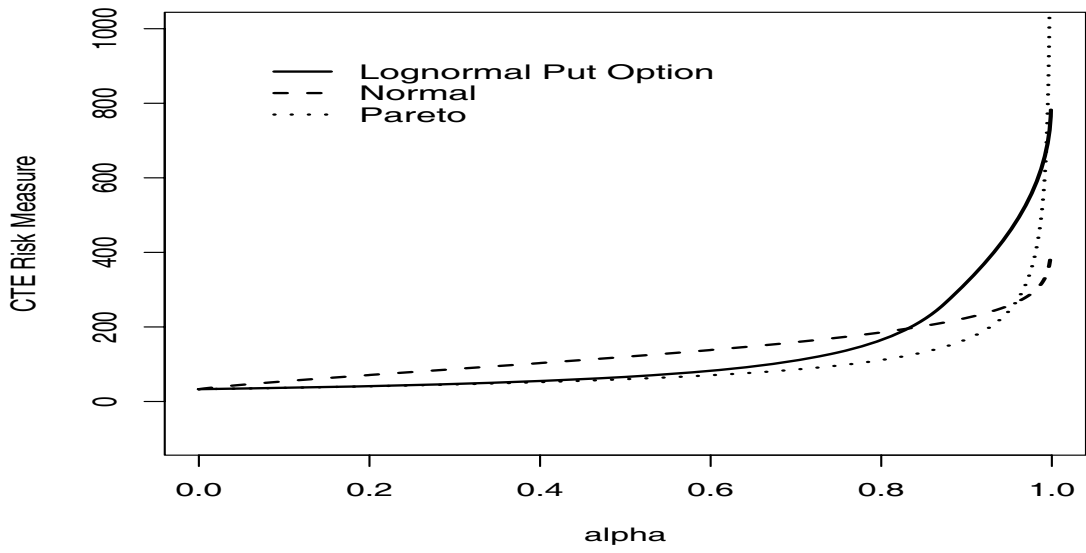



Figure 2:

The distortion function defining the α -quantile risk measure is

$$g(S(x)) = \begin{cases} 0 & \text{if } 0 \leq S(x) \leq 1 - \alpha \\ 1 & \text{if } 1 - \alpha < S(x) \leq 1 \end{cases} \quad (11)$$

The distortion function defining the α -CTE risk measure is

$$g(S(x)) = \begin{cases} S(x)/(1 - \alpha) & \text{if } 0 \leq S(x) \leq 1 - \alpha \\ 1 & \text{if } 1 - \alpha < S(x) \leq 1 \end{cases} \quad (12)$$

Note that the definition for the CTE for non-negative losses using the distortion function automatically allows for probability masses.

Exercise: For non-negative loss distributions, show that the distortion functions above are the same as the definitions in sections 2.2 and 2.3.

Other distortion risk measures include **The proportional hazard (PH) transform** (Wang (1995, 1996)):

$$g(S(x)) = (S(x))^{1/\kappa} \quad \text{for } \kappa \geq 1 \quad (13)$$

The parameter κ is a measure of risk aversion – higher values of κ correspond to a higher security level.

Example: Assume $L \sim \text{Pareto}(\gamma, \theta)$. Then the survival function is

$$S(x) = \left(\frac{\theta}{\theta + x} \right)^\gamma$$

so the distorted survival function is

$$g(S(x)) = \left(\frac{\theta}{\theta + x} \right)^{\gamma/\kappa}$$

This is a new survival function for a Pareto $(\gamma/\kappa, \theta)$ distribution. The integral of the survival function (for a non-negative loss) is the mean loss, so

$$\mathcal{H}(L) = \frac{\theta}{\frac{\gamma}{\kappa} - 1} \quad (14)$$

provided $\frac{\gamma}{\kappa} > 1$; otherwise it is undefined.

Suppose $\theta = 1200$ and $\gamma = 13$. The mean loss is \$100; the standard deviation is \$109,

and the 95% VaR is \$311. The PH-transform risk measure for, say, $\kappa = 3$ is \$360.

One problem with the proportional hazard distortion is that there is no easy interpretation of κ .

The dual power transform is defined using the distortion

$$g(S(x)) = 1 - (1 - S(x))^\kappa. \quad (15)$$

This measure has an interpretation— for integer κ , it is the expected value of the maximum of a sample of κ observations of the loss L . Suppose $\kappa = 20$, and we apply the dual power to the put option loss example above. We cannot do this analytically, but we can do a numerical integration. The result is $\mathcal{H}(X) = \$363$.

Exercise: Use Excel to calculate risk measures for the put option example,

(i) Using the PH transform with $\kappa = 20$

(ii) Using the dual power transform with $\kappa = 40$.

Answer: (i) \$745 (ii) \$479.

Wang's Transform (WT) Wang (2002) describes a shifted Normal premium principle, also in terms of a distortion function,

$$g(S(x)) = \Phi\left(\Phi^{-1}(S(x)) + \kappa\right) \quad (16)$$

Suppose, for example, we have a lognormal loss with parameters $\mu = 0$ and $\sigma = 1$. The mean loss is 1.65, standard deviation 2.16. Consider a right tail loss probability— for example, $Pr[L > 12]$. The true probability is $1 - \Phi((\log(12) - \mu)/(\sigma)) = 1 - \Phi(2.4849) = 0.0065$

The probability assigned by Wang's transform, with $\kappa = 1$, say, is found by first taking $\Phi^{-1}(0.0065) = -2.4849$. Now shift by $\kappa = 1$ to give -1.4849 . Now reapply the normal distribution function to give the distorted tail probability

$$g(S(12)) = \Phi(-1.4849) = 0.06879.$$

So the distorted probability of being in this far tail is more than ten times greater than the true probability, giving more weight to the tail in the expected value calculation. This premium principle works well with lognormal losses, as the distorted survival function is

the survival function for a lognormal distribution with shifted μ parameter.

Exercise: Show that applying Wang's Transform distortion function to a lognormal(μ, σ) loss distribution gives a lognormal($\mu + \kappa\sigma, \sigma$) distribution survival function.

Solution:

$$\begin{aligned}
 S(x) &= 1 - \Phi\left(\frac{\log(x) - \mu}{\sigma}\right) \\
 g(S(x)) &= \Phi\left(\Phi^{-1}\left(1 - \Phi\left(\frac{\log(x) - \mu}{\sigma}\right)\right) + \kappa\right) \\
 &= \Phi\left(\Phi^{-1}\left(\Phi\left(\frac{-\log(x) + \mu}{\sigma}\right)\right) + \kappa\right) \\
 &= \Phi\left(\frac{-\log(x) + \mu}{\sigma} + \kappa\right) \\
 &= 1 - \Phi\left(\frac{\log(x) - \mu - \kappa\sigma}{\sigma}\right)
 \end{aligned}$$

which is the distortion for the Lognormal ($\mu + \kappa\sigma, \sigma$), as required.

3 Coherence

With all these risk measures to choose from it is useful to have some way of determining whether each is equally useful. In Artzner et al (1997, 1999) some axioms were defined that were considered desirable characteristics for a risk measure (in fact the precursor was the discussion on 'desirable characteristics' in Gerber (1979)).

The axioms are:

Translation Invariance (TI) For any non-random c

$$\mathcal{H}(X + c) = \mathcal{H}(X) + c \tag{17}$$

This means that adding a constant amount (positive or negative) to a risk adds the same amount to the risk measure. It also implies that the risk measure for a non-random loss, with known value c , say, is just the amount of the loss c .

Positive Homogeneity (PH) For any non-random $\lambda > 0$

$$\mathcal{H}(\lambda X) = \lambda \mathcal{H}(X) \tag{18}$$

This axiom implies that changing the units of loss does not change the risk measure.

Subadditivity For any two random losses X and Y ,

$$\mathcal{H}(X + Y) \leq \mathcal{H}(X) + \mathcal{H}(Y) \tag{19}$$

The subadditivity axiom has probably been the most examined of the axioms of coherence. The essential idea is intuitive; it should not be possible to reduce the economic capital required (or the appropriate premium) for a risk by splitting it into constituent parts. Or, in other words, diversification (ie consolidating risks) cannot make the risk greater, but it might make the risk smaller if the risks are less than perfectly correlated.

Monotonicity If $\Pr[X \leq Y] = 1$ then $\mathcal{H}(X) \leq \mathcal{H}(Y)$.

Again, an intuitively appealing axiom, that if one risk is always bigger than another, the risk measures should be similarly ordered. This axiom, together with the TI axiom also requires that the risk measure must be no less than the minimum loss, and no greater than the maximum loss. This is easily shown by replacing X by $\min Y$ for the lower bound, then $\Pr[\min Y \leq Y] = 1$, so

$$\mathcal{H}(\min Y) \leq \mathcal{H}(Y) \tag{20}$$

and using axiom TI, noting that $\min Y$ is not random, $\mathcal{H}(\min Y) = \min Y$ giving the lower bound.

A risk measure satisfying the above four conditions is said to be **coherent**. The quantile risk measure is not coherent, as it is not subadditive. We can illustrate this with a simple example. Suppose we have losses X and Y , both dependent on the same underlying Uniform(0,1) random variable U as follows.

$$X = \begin{cases} 1000 & \text{if } U \leq 0.04 \\ 0 & U > 0.04 \end{cases}$$

$$Y = \begin{cases} 0 & \text{if } U \leq 0.96 \\ 1000 & U > 0.96 \end{cases}$$

Let $\mathcal{H}(X)$ denote the 95% quantile risk measure. Then

$$\mathcal{H}(X) = \mathcal{H}(Y) = 0$$

since in both cases, the probability of a non-zero loss is less than 5%. On the other hand, the sum, $X + Y$ has an 8% chance of a non-zero loss, and the 95% quantile risk measure of the sum is therefore

$$\mathcal{H}(X + Y) = 1000 > \mathcal{H}(X) + \mathcal{H}(Y) = 0.$$

The failure of quantile risk measures to satisfy the coherence axioms is one of the reasons why it is less popular with actuaries than the CTE risk measure. However, it is still in widespread use in the banking industry, where the Basel Committee on Banking Supervision introduced a 99% Value at Risk requirement, based on a 10-day trading horizon. The issues surrounding Value at Risk in banking are covered extensively in Jorion (2000)

The subadditivity axiom, and the consequent attempt to reject Value at Risk in favour of coherent measures, has been attacked on two grounds. The first is that the quantile risk measure is useful and well understood, that in most practical circumstances it is sub-additive, and the failure to be sub-additive in a few situations is not sufficiently important to reject the quantile risk measure (see, for example Danielsson et al (2005))

The second is that in some circumstances it **may** be valuable to dis-aggregate risks – though it would need to be more than a paper exercise within the company.

Exercises:

1. Show that $\mathcal{H}(X) = E[X]$ is coherent.
2. Show that the variance premium principle is not coherent.
3. The exponential premium principle for a loss X , is

$$\mathcal{H}(X) = \frac{\log\left(E\left[e^{\alpha X}\right]\right)}{\alpha} \quad \alpha > 0$$

Show that it is not coherent. Which (if any) of the coherency axioms does it satisfy?

4. Does the quantile risk measure satisfy the axioms of coherence other than subadditivity?

4 Other measures of risk

All the measures of risk listed so far are solvency measures – that is, they can be interpreted as premiums or capital requirements for financial and insurance risks. Another class of risk measures are measures of variability.

In Mean Variance Portfolio Theory, the portfolio variance or standard deviation is taken as a measure of risk. Clearly these are the most common variability measures. We assume a higher variance indicates a more risky loss random variable.

Another variability risk measure is the semi-variance measure. The motivation is that only variance on the worst side is important in risk measurement. So, instead of measuring the variance σ^2 as

$$\sigma^2 = E[(X - \mu_X)^2]$$

we only look at the worst side of the mean. Since we are dealing with a loss random variable, this corresponds to higher values of X , so the semi-variance is:

$$\sigma_{SV}^2 = E[(\max(0, X - \mu_X))^2] \tag{21}$$

which would generally be estimated by the sample semi-variance, for a sample size n ,

$$sv^2 = \sum_{i=1}^n \frac{(\max(x_i - \bar{x}, 0))^2}{n} \tag{22}$$

where $\bar{x} = \sum_{i=1}^n x_i/n$.

The mean may be replaced in the calculation by an arbitrary, known threshold parameter value, τ . This is sometimes called the threshold semi-variance, that is sv_τ^2 :

$$sv_\tau^2 = \sum_{i=1}^n \frac{(\max(x_i - \tau, 0))^2}{n} \tag{23}$$

For a profit and loss random variable Y , where a positive value indicates a profit and a

negative value a loss, the semi-variance is known as the downside semi-variance, so the threshold downside semi-variance, for example, is

$$\sigma_{\text{SV}}^2 = \text{E}[(\min(0, Y - \tau))^2] \tag{24}$$

and $\tau = 0$ would be a common threshold; that is, measuring the variance of losses with no contribution from profits.

The downside semi-deviation is the square root of the downside semi-variance.

We can use the variability risk measures to construct solvency risk measures. If we let $\nu(X)$ denote a variability risk measure, then we may construct a solvency risk measure using, for example, $\mathcal{H}(X) = \text{E}[X] + \alpha \nu(X)$. We have already used this format in the variance and standard deviation premium principles. However, none of the solvency risk measures of this form are coherent. The standard deviation principle $\mathcal{H}(X) = \text{E}[X] + \alpha \sigma$ does not satisfy the monotonicity axiom, and neither the variance principle $\mathcal{H}(X) = \text{E}[X] + \alpha \sigma^2$ nor the semi-variance principle $\mathcal{H}(X) = \text{E}[X] + \alpha \sigma_{\text{SV}}^2$ satisfy the sub-additivity axiom.

Exercise: 1. (i) *Calculate the semi-variance for the following sample of loss data:*

(ii) *Calculate the threshold semi-variance, with a threshold of $\tau = 35$.*

1, 1, 1, 2, 5, 8, 35, 75

Exercise: *Show that the standard deviation premium principle is not coherent. Which (if any) of the coherency axioms does it satisfy?*

Exercise: *Show that the variance premium principle is not coherent. Which (if any) of the coherency axioms does it satisfy?*

5 Estimating risk measures using Monte Carlo simulation

5.1 The Quantile Risk Measure

In actuarial applications we often use Monte Carlo simulation to estimate loss distributions, particularly when the underlying processes are too complex for analytic manipulation.

Using standard Monte Carlo simulation² we generate a large number of independent simulations of the loss random variable L . Suppose we generate N such values, and order them from smallest to largest, such that $L_{(j)}$ is the j -th smallest simulated value of L . We assume the ‘empirical’ distribution of $L_{(j)}$ is an estimate of the true underlying distribution of L .

For example, suppose we use Monte Carlo simulation to generate a sample of 1000 values of a loss random variable. We are interested in the 95% quantile risk measure and the 95% CTE for the loss. We have two important questions; one is, how do we use the output to estimate the risk measures; the second is, how much uncertainty is there in the estimates?

Suppose we want to estimate the 95%-quantile risk measure of L . An obvious estimator is $L_{(950)}$. It’s an obvious candidate because it’s the 95% quantile of the empirical distribution defined by the Monte Carlo sample $\{L_{(j)}\}$ – that is 95% of the simulated values of L are less than or equal to $L_{(950)}$. On the other hand, 5% of the sample is greater than or equal to $L_{(951)}$, so that’s another possible estimator. In *Loss Models* (Klugman Panjer and Willmot (2004)), the ‘smoothed empirical estimate’ suggested is to assume that $L_{(j)}$ is an estimate of the $j/(N + 1)$ quantile of the distribution, and use linear interpolation to get the desired quantile. That means that $L_{(950)}$ is assumed to be an estimate of the $950/1001 = 94.905\%$ quantile, and $L_{(951)}$ is assumed to be an estimate of the $951/1001 = 95.005\%$ quantile. Linear interpolation for the 95% quantile would give

$$Q_{95\%} \approx (0.05) L_{(950)} + (0.95) L_{(951)}$$

Generalizing, we have three possible estimators for Q_α , from a sample of N simulated

²We assume no variance reduction techniques.

values of a loss random variable L , where $N\alpha$ is assumed to be an integer.

1. $L_{(N\alpha)}$
2. $L_{(N\alpha+1)}$
3. Interpolate between $L_{(N\alpha)}$ and $L_{(N\alpha+1)}$, assuming $L_{(r)}$ is an estimate of the $r/(N+1)$ quantile – the ‘*smoothed empirical estimate*’.

None of these estimators is guaranteed to be better than the others. Each is likely to be biased, though the bias is generally small for large samples. We cannot even be certain that the true α -quantile lies between $E[L_{(N\alpha)}]$ and $E[L_{(N\alpha+1)}]$. In practice, for the actuarial loss distributions in common use, where we are interested in the right tail of the loss distribution, we usually get lower bias from using either $L_{(N\alpha+1)}$ or the smoothed empirical estimate. It is also worth noting that all three estimators are asymptotically unbiased, so for large samples the bias will be small. Also, the bias will tend to be larger in the farther tails of the distribution.

Table 1 is an excerpt from one sample of 1000 values simulated from the example Normal loss distribution, with a mean of 33 and a standard deviation of 109.0. These are the 100 largest values simulated:

Exercise: Use the information in Table 1 to estimate the 95th and 99th quantile for the Normal (33, 109²) loss, using the three estimators above. What is the relative ‘error’ – that is, the difference between your values and the true values from Section 1, expressed as a percent of the true values?

Of course, in practice, when we use Monte Carlo, we do not know the true quantile values. Suppose we use $L_{(N\alpha)}$ as an estimate of the α -quantile. This estimate will be subject to sampling variability. We can use the simulations around the estimate to construct a non-parametric confidence interval for the true α -quantile, Q_α for the distribution³.

The number of simulated values falling below the true α -quantile, Q_α is a random variable, M , say, with a binomial distribution. Each simulated value of L either does fall below Q_α – with probability α – or does not, with probability $(1-\alpha)$. So

$$M \sim \text{binomial}(N, \alpha)$$

³Strictly, for the $N\alpha$ order statistic

which means that

$$E[M] = N\alpha \quad V[M] = N\alpha(1 - \alpha)$$

Suppose we want a 90% confidence interval for Q_α . We first construct an 90% confidence interval for $E[M]$, say, (m_L, m_U) such that

$$\Pr[m_L < M \leq m_U] = 0.9$$

and we constrain the interval to be symmetric around $N\alpha$, so $m_L = N\alpha - a$, say, and $m_U = N\alpha + a$. So, if $F_M(x)$ is the binomial distribution function for M

$$F_M(N\alpha + a) - F_M(N\alpha - a) = 0.9$$

using the Normal approximation to the binomial distribution gives:

$$a = \Phi^{-1}\left(\frac{1 + 0.9}{2}\right) \sqrt{N\alpha(1 - \alpha)} \tag{25}$$

The 90% confidence interval for $E[M]$ gives the range of ordered Monte Carlo simulated values corresponding to a 90% confidence interval for Q_α :

$$\Pr\left[Q_\alpha \in \left(L_{(N\alpha-a)}, L_{N\alpha+a}\right)\right] = 0.9$$

In practice, if a is not an integer, we would round up to the next integer, although interpolation would also be acceptable.

The $L_{(950)}$ estimate from this simulated sample is 209.2. Suppose we want to construct a 90% confidence interval for this estimate. We find

$$a = (\Phi^{-1}(0.95))\sqrt{1000(0.95)(0.05)} = (1.645)(6.892) = 11.33$$

Round up to 12, then we have a 95% confidence interval for $Q_{95\%}$

$$(L_{(938)}, L_{(962)}) = (200.5, 231.4)$$

So, in general, the process for the non-parametric q -confidence interval for the $N\alpha$ order statistic is

1. Calculate

$$a = \Phi^{-1}\left(\frac{1 + q}{2}\right) \sqrt{N\alpha(1 - \alpha)}$$

2. Round a up to the next integer
3. The q -confidence interval is $(L_{(N\alpha-a)}, L_{(N\alpha+a)})$.

For the Normal approximation to be valid, we need $N(1 - \alpha)$ (or $N\alpha$ if smaller) to be at least around 5.

Exercise: Use the non-parametric method, and the information in Table 1, to estimate the 95% confidence interval for $Q_{92.5\%}$ in the Normal example.

Solution: (174.3, 203.7)

Another way to explore the uncertainty would be to repeat the simulations many times. By this we mean simulate a large number, R , say, of samples, each of N values. Each simulated sample can be used to estimate the quantile risk measure; let $\hat{Q}_\alpha(i)$ denote the estimate from the i -th simulated sample. The R values of $\hat{Q}_\alpha(i)$ then can be viewed as an i.i.d. sample. Then we would use the mean of these as the estimated risk measure, that is:

$$\hat{Q}_\alpha = \frac{1}{R} \sum_{i=1}^R \hat{Q}_\alpha(i)$$

and the sample standard deviation of the risk measures is an estimate of the standard error:

$$\hat{s}_Q = \frac{1}{R-1} \sum_{i=1}^R (\hat{Q}_\alpha(i) - \hat{Q}_\alpha)^2$$

Then we can use the sample standard deviation to construct an approximate confidence interval for the risk measure; for example, for a 90% confidence interval we would use

$$(\hat{Q}_\alpha - 1.64\hat{s}_Q, \hat{Q}_\alpha + 1.64\hat{s}_Q)$$

We can illustrate this with the Normal(33, 109²) example. With 1000 replications, each with a sample size of 1000, we find the mean values and standard deviations of the Monte Carlo estimators are:

$$\bar{L}_{(950)} = 211.53 \quad s_{L_{(950)}} = 7.30$$

$$\bar{L}_{(951)} = 212.61 \quad s_{L_{(951)}} = 7.34$$

Mean Smoothed Estimator = 212.56

The true 95% quantile value is 212.29; the smoothed estimator has a relative error averaging 0.13%. Using $L_{(950)}$ gives a relative error averaging -0.36%, and using $L_{(951)}$ gives a relative error averaging 0.15%.

Similarly, the 99% estimators, this time using 5000 replications of the sample, are:

$$\bar{L}_{(990)} = 284.41 \quad s_{L_{(990)}} = 12.5$$

$$\bar{L}_{(991)} = 288.41 \quad s_{L_{(990)}} = 12.9$$

Mean Smoothed Estimator = 288.37

which compares with the true 99% quantile value of 286.57. We note that nearer the tail the standard deviation and the relative errors increase.

Using $L_{(951)}$ as the estimator for $Q_{95\%}$, the mean value is 212.61, and the associated standard deviation from our 1000 simulations is $s_{Q_{95\%}} = 7.34$. This gives a 90% confidence interval for $Q_{95\%}$ of (200.57,224.65), which is similar to the non-parametric 90% confidence interval, but achieved at a much greater cost, of an additional 999 simulated samples, each with 1000 simulated loss values.

5.2 CTE

The CTE is the mean of the worst $100(1 - \alpha)\%$ of the loss distribution, so we estimate the CTE using the mean of the worst $100(1 - \alpha)\%$ simulations, that is, assuming $N(1 - \alpha)$ is an integer,

$$\widehat{\text{CTE}}_{\alpha} = \frac{1}{N(1 - \alpha)} \sum_{j=N\alpha+1}^N L_{(j)} \quad (26)$$

For example, the figures in Table 1 show the worst 100 simulations from a sample of $N=1000$. To estimate the 95% CTE we average the worst 50 values, giving a CTE estimate of \$260.68. This compares with the true value of \$257.83, which was calculated in Section 2.3.

$L_{(901)}$ to $L_{(910)}$									
169.1	170.4	171.3	171.9	172.3	173.3	173.8	174.3	174.9	175.9
$L_{(911)}$ to $L_{(920)}$									
176.4	177.2	179.1	179.7	180.2	180.5	181.9	182.6	183.0	183.1
$L_{(921)}$ to $L_{(930)}$									
183.3	184.4	186.9	187.7	188.2	188.5	191.8	191.9	193.1	193.8
$L_{(931)}$ to $L_{(940)}$									
194.2	196.3	197.6	197.8	199.1	200.5	200.5	200.5	202.8	202.9
$L_{(941)}$ to $L_{(950)}$									
203.0	203.7	204.4	204.8	205.1	205.8	206.7	207.5	207.9	209.2
$L_{(951)}$ to $L_{(960)}$									
209.5	210.6	214.7	217.0	218.2	226.2	226.3	226.9	227.5	227.7
$L_{(961)}$ to $L_{(970)}$									
229.0	231.4	231.6	233.2	237.5	237.9	238.1	240.3	241.0	241.3
$L_{(971)}$ to $L_{(980)}$									
241.6	243.8	244.0	247.2	247.8	248.8	254.1	255.6	255.9	257.4
$L_{(981)}$ to $L_{(990)}$									
265.0	265.0	268.9	271.2	271.6	276.5	279.2	284.1	284.3	287.8
$L_{(991)}$ to $L_{(1000)}$									
287.9	298.7	301.6	305.0	313.0	323.8	334.5	343.5	350.3	359.4

Table 1: Largest 100 values from a Monte Carlo sample of 1000 values from a Normal(33, 109²) distribution

Exercise: Compare the 99% CTE estimate from the Table 1 Monte Carlo sample with the true value for the $N(33,109^2)$ distribution.

Solution: The estimate from the sample is \$321.8, compared with the true value \$323.5

The most obvious candidate for estimating the standard deviation of the CTE estimate is $s_1/\sqrt{N(1-\alpha)}$ where s_1 is the standard deviation of the worst $100(1-\alpha)\%$ simulated losses:

$$s_1 = \sqrt{\frac{1}{N(1-\alpha)-1} \sum_{i=N\alpha+1}^N (L_{(j)} - \widehat{CTE}_\alpha)^2}$$

We might use this because, in general, we know that the variance of the mean of a sample is equal to the sample variance divided by the sample size. However, this will underestimate the uncertainty, on average. The quantity that we are interested in is $V[\widehat{CTE}_\alpha]$. We can condition on the quantile estimator \hat{Q}_α , so that:

$$V[\widehat{CTE}_\alpha] = E[V[\widehat{CTE}_\alpha|\hat{Q}_\alpha]] + V[E[\widehat{CTE}_\alpha|\hat{Q}_\alpha]] \quad (27)$$

The plug-in $s_1^2/N(1-\alpha)$ estimates the first term; we need to make allowance also for the second term which considers the effect of the uncertainty in the quantile.

A way to allow for both terms in (27) is using the influence function approach of Manistre and Hancock (2005). The variance of the CTE estimate can be estimated using

$$s_{CTE_\alpha}^2 = \frac{s_1^2 + \alpha(\widehat{CTE}_\alpha - \hat{Q}_\alpha)^2}{N(1-\alpha)}$$

Using the Table 1 data, with $\alpha = 0.95$, for example, we have $\hat{Q}_{95\%} = 212.56$ using the smoothed estimator, and the 95% CTE is estimated above as \$260.68. The standard deviation of the largest 50 values in Table 1 is 37.74. Also, $N = 1000$, $\alpha = 0.95$, so the standard deviation estimate is:

$$\sqrt{\frac{37.73^2 + 0.95(260.68 - 212.56)^2}{50}} = 5.42$$

Note that the first term in this formula is the same $s_1/\sqrt{N(1-\alpha)}$ term that we proposed earlier; the second term allows for the uncertainty in the quantile.

Another approach to the standard error is to repeat the sample simulation a large number of times and calculate the standard deviation of the estimator, exactly as we did for the quantile. This is effective, but expensive in terms of the volume of additional simulation required.

Exercise: Estimate the 99% CTE and its standard error using the data in Table 1.

6 More Exercises

1. Assume an exponential loss distribution with mean θ . Derive an expression for the difference

$$\text{CTE}_\alpha - Q_\alpha$$

2. For a loss random variable L , derive a relationship between the α -CTE of L and the mean residual lifetime function (Klugman, Panjer and Willmot).
3. Assume a Weibull loss distribution with $\theta = 1000$, $\tau = 2$ (using the notation of KPW).
 - (a) Calculate the 95% quantile risk measure.
 - (b) Calculate the 95% CTE (Use the limited expected value function; you will need the gamma and incomplete gamma functions)
 - (c) Calculate the distortion risk measure using Wang's PH transform, $\kappa = 5$.
4. You are given the following discrete loss distribution:

$$X = \begin{cases} 10 & \text{with probability } 0.8 \\ 50 & \text{with probability } 0.1 \\ 200 & \text{with probability } 0.08 \\ 1000 & \text{with probability } 0.02 \end{cases}$$

Calculate:

- (a) The 95% quantile
- (b) the 95% CTE

- (c) The semi-variance
- (d) The distortion risk measure using Wang's second premium principle, with $\kappa=0.5$.

References

- Artzner, P., Delbaen F., Eber, J.M., and Heath, D. (1997). Thinking Coherently. *Risk*, 10, November, 68-71.
- Artzner, P., Delbaen F., Eber, J.M., and Heath, D. (1999). Coherent Measures of Risk. *Mathematical Finance*, **9**, 203-228.
- Bühlmann (1970) *Mathematical Methods in Risk Theory* Springer-Verlag Berlin.
- Danielsson J., Jorgenson B. N, Samorodnitsky G., Sarma M., de Vries C. G. (2005) Sub-additivity re-examined: the case for Value at Risk. Working Paper available from www.RiskResearch.org
- Gerber H. U. (1979) *An Introduction To Mathematical Risk Theory*. S.S. Huebner Foundation, Wharton School, University of Pennsylvania. Distributed by R Irwin, Philadelphia.
- Klugman S., Panjer H. and Willmot G. (2004) *Loss Models: From data to decisions*. (2nd Ed.) Wiley.
- Jorion, P. (2000). *Value at Risk (2nd Edition)*. McGraw-Hill.
- Manistre J. and Hancock G. (2005) Variance of the CTE Estimator. *North American Actuarial Journal*, **9**, 2, 129-136
- Wang, S. (1995). Insurance Pricing and Increased Limits Ratemaking by Proportional Hazards Transforms. *Insurance: Mathematics and Economics* **17** 43-54.
- Wang, S. (1996). Premium Calculation by Transforming the Layer Premium Density. *ASTIN Bulletin* **26** 71-92.
- Wang, S. (2002), A universal framework for pricing financial and insurance risks . *ASTIN Bulletin* **32**, 213-234.