

## Text Lecture 3.2 – Special VaRs and the Expected Shortfall

Hi there, welcome to this second class.

Today we will start with an exercise, because it is very important that you are familiar with the computation of the VaR.

So let's read the text of the problem. A 1-year project has a 94% probability of leading to a gain of 5 million. A 3% change of a gain of 2 million euros. And a 2% change of leading to a loss of 3 million. And a 1% chance of producing a loss of 8 million. The question is: what is the VaR at alpha level 0.98...so the 98% VaR? And what happens if alpha is 0.99?

To solve this type of exercises, it may be convenient (at least at the very beginning) to write all the data we have in table format.

Our table starts with three columns: losses, probability as a percentage and probability as a decimal.

Losses must be sorted, from the smallest to the largest. We can use quantities in million.

A gain of 5 corresponds to a loss of -5. The same for a gain of 2. Then we have 3 and 8 for actual losses.

To every loss, we can associate the corresponding probability, as a percentage and as a decimal.

We then add a fourth column, containing the cumulative probabilities as decimals: 0.94, 0.97, 0.99 and 1.

Thanks to this table, we can now draw the empirical cumulative distribution function of losses. On the x axis we put losses. On the y axis the cumulative probabilities. For reasons of space, the scale of the y axis is not really representative here.

Now, we start with the smallest loss: -5. The corresponding cumulative probability is 0.94. This means that we can draw a vertical segment, going from the x axis in -5 to the point (-5, 0.94).

Then we know that the cumulative probability does not change until we reach -2, where it increases to 0.97. Hence we first have a horizontal segment and then a vertical one.

And so on, so that we can draw the entire empirical cumulative distribution function (also known as ecdf).

Now, the first question asks us to find the 1-year 0.98 VaR.

The trick is really simple.

Let's consider the y-axis. Is 0.98 there? No. But we can add it between 0.97 and 0.99.

Now, let's draw a dotted horizontal line passing through 0.98. In which point does this line cross the empirical cumulative distribution function? Here.

And this point corresponds to a loss of 3. This is our 0.98 VaR! Fine.

For the 0.99 VaR we do the same, and 0.99 is already in our plot. Good!

Let's draw the horizontal line.

Wait a minute. The line does not cross the ecdf. It overlaps for an entire segment, from loss 3 to loss 8. So, what's the VaR?

Here we can use a convention. In situations like this, the VaR corresponds to the average loss in the segment. This means that we compute  $8+3=11$  and then we divide by 2. The 99% VaR is 5.5.

Good, we have solved our problem.

Another type of VaR, which is often used is the so-called mean-VaR, that is to say, the VaR centered around the mean of the loss distribution. It is simply given by the difference between the VaR and the mean  $\mu$ . As we will see later on, the mean-VaR is used to determine the economic capital against losses in loans.

If we assume a specific distribution for the loss distribution, we can obtain special formulas for the VaR. For example, for a Gaussian distribution, the VaR is simply computable using the quantile function of a standard Gaussian, that is to say a  $\text{Normal}(0,1)$ . The VaR alpha for a normal distribution, with mean  $\mu$  and standard deviation  $\sigma$ , is equal to  $\mu + \sigma$  times the alpha quantile of a standard Gaussian. For the mean VaR, naturally the  $\mu$  term disappears.

And this is the formula, if the loss distribution is a t-distribution.

Please notice the notation: "var" in small letters is the variance and not the Value-at-Risk!

And now, another exercise. We know that the historical 1-year distribution for our portfolio of loans is well approximated by a  $\text{Normal}(10,5)$ . That is to say a normal distribution with mean  $\mu=10$  and standard deviation  $\sigma=5$ .

Now, the question is: what are the 95% and 98% VaRs?

To solve this exercise we can use the standard normal tables or R.

The 0.95 quantile of a standard Gaussian is 1.6648. This is the threshold value below which 95% of all observations lie according to a standard Gaussian.

The 98% quantile is 2.0537. These quantiles can be computed in R using the norm function. If you just type `qnorm(0.95)`, you get the quantile of the standard Gaussian. However, notice that you can also write `qnorm(0.95,10,5)`.

In this case the quantile is directly computed from a `Normal(10,5)`.

18.2243 is the value of our 95% VaR. We can obtain this directly from `qnorm`, or using the formula in slide 4. For the 98% VaR the reasoning is exactly the same.

Now assume that things go bad, and that we can observe a loss which is greater than our VaR alpha. Now, a natural question we may want to answer is the following: what's the expected loss?

The expected shortfall is the statistical quantity that tries to answer this question. From a statistical point of view, the expected shortfall at level alpha is a sort of mean excess function, i.e. the average value of all the values exceeding a special threshold, the VaR alpha!

Why is the expected shortfall important? Simple: two loss distributions may have the same value of VaR alpha, but two different expected shortfalls. This is due to the fact that the shape of the right tail may be different. For example, if we compare the two pictures, we can expect the expected shortfall of the right-hand side distribution to be larger, given that this distribution assigns a higher probability to larger losses' values.

Let's consider a last exercise to understand how we can compute the expected shortfall.

A portfolio of loans may lead to the losses in the table. What is the expected shortfall for alpha equal to 0.95? And 0.99? By the way, what is the 95% VaR in this case? Can you see it immediately?

To compute the expected shortfall, the trick is to sort the losses as usual. Then we start from the bottom, i.e. from the largest losses and we go backward, summing the corresponding probabilities, until we reach the 1-alpha level. If alpha is 95% then 1-alpha is 5%. Look at the table: the 25 million loss has a chance of 0.5%. Is this equal to 5%? No. Hence we sum the probability of the 20 million loss, getting 3%. Is this equal to 5%? Not yet, therefore we also sum the probability of the 12 million loss. Now the total is  $0.5+2.5+2=5$ .

Here we are. Summing the three largest losses the total probability is 5%. The 95% expected shortfall is nothing more than the weighted average of these three losses, where the weights are their occurrence probabilities, and the denominator is 5%. Our expected shortfall is therefore 17.3 million.

In the 99% case, things are a little bit different. If we sum 0.5% (the probability of a loss of 25 million) with 2.5% (for a loss of 20 million), the result is greater than 1%, our 1 minus alpha. So what?

In this case we only have to consider the part of the 20 million loss probability that, together with the initial 0.5% of the 25 million loss, sums up to 1%.

This quantity is clearly another 0.5%. Hence the 99% ES is computed as: 20 times 0.005, plus 25 times 0.005, over 0.01. The result is 22.5 million.

As you can imagine, even for the expected shortfall we can obtain special closed formulas, depending on the theoretical loss distribution we assume.

This is what we get for a Gaussian distribution.

Ok. This has been a rather long class. It is time to stop.

See you next time. Bye.