Text Lecture 3.1 – Introducing Value-at-Risk

Welcome everybody. This week we will consider a fundamental but somehow controversial measure of risk: the VaR, or Value-at-Risk.

The VaR is a measure that tries to answer a simple but significant question: How bad can things get, in terms of losses, when we invest, we lend money, and so on?

In more probabilistic terms, we look for a measure that tells us: With probability alpha we will not lose more than V euros (or dollars, or pounds, and so on) in time T.

The quantity V is the VaR. Alpha is the so-called confidence level. While capital T is the time horizon over which the VaR is computed.

The VaR can be computed using two different distributions: the distribution of gains, where a loss is a negative gain. Or the distribution of losses, where a gain is a negative loss.

We will prefer this second approach, but nothing really changes from a conceptual point of view. From a practical point of view, we could say that the big difference is a change of sign for the VaR, because of symmetry considerations.

Using some more formality, the VaR is nothing but a quantile of the loss distribution, and in particular the alpha-quantile for which the probability of observing a larger loss, given the available information, is equal to 1-alpha.

In other terms, the VaR alpha is the threshold loss such that the probability of observing smaller losses, given the loss distribution, is alpha.

Given a loss distribution, the 90% VaR is the threshold loss for which 90% of losses are smaller and only 10% are larger. Naturally, always with respect to the loss distribution under scrutiny.

The equation you can see on your screen is sometimes represented with a greater than/equal to sign. As you can imagine this makes no difference for continuous loss distributions, as typically assumed in credit risk models.

The quantity capital F is the cumulative distribution function of losses.

If you are not familiar with this terminology and concepts, such as for example cumulative distribution and quantile, please refer to the prerequisites of this course.

In the field of Credit Risk, the VaR is often referred to as C-VaR (credit VaR), as we will see later in this course. For the moment, VaR and C-VaR are just synonyms for us.

The VaR plays a major role in market and operational risks too.

The Value-at-Risk essentially depends on 2 elements: the loss distribution, and the alpha value. A loss distribution is always expressed over a time horizon T and it can be empirical or theoretical.

In the first case, it is the so-called historical distribution, that is the distribution that emerges from the observation of reality, when we collect data about historical losses.

In the second case, it can be whatever distribution and it is essentially used for modeling purposes.

In week 6 we will see more about this.

As far as the choice of the time horizon, capital T, is concerned, it should reflect the time period over which a financial institution is committed to hold its portfolio.

This period may be affected by contractual and legal constraints and/or liquidity considerations.

In the case of credit risk, these constraints are given by the rules of Basel II and III.

The alpha value, from a theoretical point of view, may be freely chosen by the risk manager. In reality, it is often determined by law or other prescriptions. Common values are 0.95, 0.99, 0.995 and 0.999.

For Credit Risk, we are usually interested in the VaRs you see on screen.

The most used alpha levels are 99% and 99.9%. The most frequent time horizons are 1 day, 10 days and 1 year.

In particular, the 1-year VaR at 99.9% confidence level is fundamental when defining the capital requirements for the banking book. Again, we will come back to these topics later on during the course.

Ok, let's now consider the present exercise. Suppose that, for a 1-year project, all the outcomes between a gain of 80 million and a loss of 20 million are considered equally likely. What is the VaR for=0.9, that is to say at the 90% confidence level?

The text of the problem gives us an important information: all the outcomes between a gain of 80 million and a loss of 20 million are considered equally likely.

This means that our loss distribution is represented by a uniform distribution over the support [-80,20]. Since we have no additional information, we may assume that this distribution is continuous.

Hence we have a continuous uniform with support [-80,20].

Let us make a simple drawing. On the x axis we put losses. So, what is the 90% VaR?

In other words, what is the threshold loss such that the probability of observing a larger loss in only 10%?

Clearly, given the uniform distribution of losses between -80 and +20, this quantity is 10 million euros. Can you see this?!

See you next time. Bye!