TW3421x - Week 7 - Stress Testing

Hi there, if you have read financial newspapers in the last months, you have probably read about stress testing.

What is stress testing?

You have probably read about the fact that now, under Basel III, banks are required to perform stress testing, in order to take into account the risk of extreme events.

And this is true for market risk, and credit risk.

So, the aim of this class is to introduce you to stress testing.

A stress test is an procedure, a way to determine the ability of a given financial institution, typically a bank, to cope with an economic crisis, or bad economic conditions, in general.

Using stress testing, a bank must answer questions like:

what happens if interest rates increase by at least x%?

what happens if the correlation among defaults increases?

And so on.

Thanks to Basel III, stress testing has become increasingly important.

It is now a regulatory requirement for large international banks, which must prove to have adequate capital allocations to cover potential losses due to extreme, but plausible, events.

A stress testing procedure is usually based on the so-called scenarios.

A scenario is simply a given configuration of parameters and variables for the model we use in assessing risk. Credit risk in our case.

The idea is simply to see what happens for unusual values, a scenario, of macroeconomic quantities such as interest rates, inflation rates, unemployment rates, volatility, etc.

We want to see what happens to our risk measures and capital requirements under particularly stressed conditions.

From an operative point of view, stress testing is performed using computational and statistical tools, such as Monte Carlo simulations, extreme value theory and, as the statisticians among you may guess, sensitivity analysis.

We can use a simple example to give an idea of how stress testing works.

Assume that for estimating the PD of a counterparty, and as a consequence the capital requirements under the F-IRB approach, we use Merton's model, so that the 1-year probability of default of our counterparty is given on your screen.

Check Week 5 for more details.

Now assume that r=0.02, V0=22, B=16, σV=0.2. Then the 1-year probability of default is 5.57%.

A simple scenario is to assume that sigmaV increases to 0.5, or 0.8, or even 1.0 and see what happens to the PD.

Using the three values of sigmaV, we see that the probability of default of our counterparty dramatically increases. All this is likely to make our capital requirements increase as well.

Do we have enough capital for that?

A bank decides if it is ready to cope with such extreme cases, by computing the capital requirements under different scenarios.

The regulator usually decides which scenarios need to be taken into consideration, when computing capital allocations.

The previous example is just a trivial one.

Naturally we can also think of more complex scenarios, in which more than one single element changes.

This is what banks do using complex computer programs.

Scenarios may be produced on the basis of:

Historical evidence. Expert judgments. Decisions of the regulator.

Historical evidence means that we want to test how we are able to cope with extreme scenarios we have already observed in the past.

When dealing with expert judgments, we want to test hypotheses developed by economists about future trends and cycles.

Finally, there are cases in which a specific scenario may be imposed by the regulator. This is something that changes from country to country. For example it has become rather common in Europe, where the ECB really insists on stress testing.

Apart from scenario analysis, another way of performing stress testing is to use the so-called stressed measures of risk.

These are the same measures of risk we have seen together, but we use them in more "extreme" situations.

A simple example is the Stressed VaR.

Developed for Market Risk, Stressed VaR, also known as S-VaR, is now popular in Credit Risk as well.

The computation of S-VaR follows the rules of VaR, but we only consider the worst losses, the worst scenarios, typically the worst 50%, that is the largest ones.

A simple example in R is the best way to understand stressed value-at-risk.

Assume we know that our losses are lognormally distributed with mean 1 and standard deviation 2.

Let's start by generating 100 losses using the function rlnorm.

This is just an example, in real life, one typically uses historical data.

As usual, we can plot a histogram to see how our losses are distributed.

The 95% VaR is easily obtained using the quantile function, as we have already seen.

The command is quantile(losses,0.95,type=3).

Type 3 means that we want to identify the closest value in the data set, without interpolations

I suggest you to type ?quantile to check again the properties of the quantile function.

Our Var is 61.48.

In order to compute the 95% S-VaR, we first sort all losses, from the smallest to the largest.

Then we select the top 50%. In our case, with 100 losses, this means the 50 largest losses.

We store these losses into a new variable.

Then we can again apply the quantile function.

The S-VaR is 155.17. Much larger!

This VaR is in fact computed on the worst losses!

Have you notice that the 95% S-VaR corresponds to a 97.5% VaR on the original losses?

Could you say why?

My suggestion is to think about 1-alpha…

Ok, see you next time.

Bye!