

Calculating and communicating tail association and the risk of extreme loss

Abstract of the London Discussion

[Institute and Faculty of Actuaries, 26 September 2011]

Contact

Paul Sweeting, F.I.A., E-mail: paul.j.sweeting@jpmorgan.com

This abstract relates to the following paper: Sweeting, P. and Fotiou, F. Calculating and communicating tail association and the risk of extreme loss. *British Actuarial Journal*, doi:10.1017/S1357321712000347

Mr S. Shepley, F.I.A. (the Chair): I am looking forward to an interesting discussion. It is a topic that I, for one, have come across many times. Quite a long time ago now, I was sitting in a board room in Bermuda with \$1.4 billion of capital with which to play. The topic was: how do we communicate to our bosses where we want to invest this capital, making use of a lot of capital modelling? That is now in vogue for everybody, and gives rise to the question of how we communicate the results of all our sophisticated capital modelling in a way that is compelling and allows people to make not just a decision but an informed decision?

I congratulate the authors of the paper for bringing this topic to the fore. If you look at the financial crisis there was a lot of capital which was effectively backing tail risk in the banking sector. Currently, there is a lot of discussion with various regulators, not just the Financial Services Authority (FSA), but the Swiss Financial Market Supervisory Authority (FINMA) and others, around tail risk and how that should be allowed for in a Solvency II or a regulatory capital model. So there are many reasons to need to communicate on this topic. I am looking forward to hearing more about how we can go about doing that. I know the authors are very keen to obtain your contributions as to how some of this might play a part in routine business, as management of capital becomes far more important in the way that businesses and executive boards take their decisions.

I should now like to turn the meeting over to Paul Sweeting and Fotis Fotiou to introduce their paper.

Mr P. Sweeting, F.I.A. (introducing the paper): Firstly, there are some acknowledgements I would like to make. The majority of this work was carried out whilst we were at the University of Kent, and it was a grant from the UK Actuarial Profession that made it possible for us to carry out this work.

I have also had useful input on the scope and content of this work from the Extreme Events Working Party and the Enterprise Risk Management Practice Executive Committee (ERM PEC) Research and Thought Leadership Sub-Committee.

The areas we will cover today can be divided in two ways. The first is the areas of tail association and of extreme loss. Measures of tail association consider the extent to which extreme values for

two or more variables are likely to occur together, something ignored if only the correlation between data series is used. This is important when considering the structure and parameterisation of financial models, particularly when those models are being used to measure risk in extreme circumstances. This is because focusing on the jointly extreme observations can help ensure that the structure of a model is sound. However, there are a number of ways in which tail association can be measured, and some approaches are better, and easier to apply, than others.

However, measures of the risk of extreme loss are also important, often more so for an investor. This is because it is extreme loss that can result in insolvency. It is important to recognise that the risk of extreme loss is not just a subset of the risk of extreme co-movement or tail association. This is because, whilst jointly extreme observations can cause a combination of large losses, large losses are not exclusively caused by jointly extreme observations. In particular, a large loss can occur with an extreme value from one variable and an average value from another, a scenario ignored if the likelihood of jointly extreme variables is the sole consideration. Joint loss also depends on the marginal distributions of the risks, that is to say the value of the observations is important. Whilst this holds for measures of co-movement based on the linear correlation coefficient, it is not generally true: most measures of tail association look only at the order of observations, as measured by the copula between them. The risk of joint loss also depends on how much of each risk you have. So, whilst a measure of tail association might only look at the fact that the relationship between, say, UK and US equity returns is significant or not, the risk of joint loss will be different depending on the allocation to the two asset classes.

The second way in which the content of this paper can be divided is into the issues of calculation and of communication. Calculation covers the derivation of quantitative measures that describe the level of tail association or the risk of extreme loss. However, you can end up with information on a large number of relationships, and it's important to be able to communicate this information efficiently. This means that showing results in a graphical format form a particularly important part of this paper.

Before handing over to Mr Fotiou, I think it is important to note that we are not recommending that particular methods of calculation or communication should be adopted dogmatically. We hope to give some ideas as a starting point, but it is for users to decide how they wish to calculate and communicate measures that are relevant to them.

Mr F. Fotiou: What I would like to talk about is the work we carried out on measures of tail association and the risk of extreme loss, and some of the conclusions to which we came. First, let us look at tail association. In broad terms, we considered two approaches for measuring tail association. The first of these was tail correlation and the second was tail dependence. Tail correlation can be thought of as describing the shape and direction of the relationship between extreme observations, whilst tail dependence is concerned more with the proportion of observations that are found in the tail.

Calculating tail correlation is actually quite straightforward. Essentially, all you are doing is working out a standard correlation coefficient for a sub-group of all observations, that is, those in the tail. You can use any sort of correlation coefficient for this purpose, although rank correlation coefficients are more robust since the linear correlation coefficient is only a valid measure of association if the data are jointly elliptical.

It is also worth noting that correlation coefficients do not have to describe the strength of a relationship between just two variables. It is possible to extend a coefficient to cover three, four or

even more variables. As with the standard bivariate measures of correlation, these multi-dimensional statistics can be turned into tail measures. However, unlike the bivariate measures there are a number of different versions of each multi-dimensional correlation statistic.

One issue with tail correlation is that you need to define what constitutes “the tail”. However, since actuaries need to deal with subjectivity all the time, this does not mean that tail correlation should not be used.

A more important issue is that tail correlation is concerned with the strength of the relationship within the tail, ignoring the fact that many extreme observations may lie off the main diagonal. This is important because it means that a strong relationship between variables in the tail does not necessarily imply a high level of risk. Mr Andrew Smith has done some interesting work here. He has defined arachnitude, a measure which is scaled by extreme observations that are spread in the four corners of a scatter plot like a spider. The relationship between two variables in the tail might be weak, but if the tail contains a large proportion of the observations, then the risk is still high.

Tail dependence overcomes this issue directly. It calculates the proportion of observations in the tail relative to the maximum proportion that there could possibly be. This is usually defined in terms of the quantiles of two distributions, u , and the copula which joins them, $C(u,u)$. The maximum proportion of observations in the bottom corner bounded by the u^{th} quantile for each of two variables is actually equal to u , so in general terms tail dependence can be measured by $C(u,u)/u$.

The most common measure of tail dependence is the coefficient of tail dependence. This is calculated, for the coefficient of lower tail dependence, by evaluating $C(u,u)/u$ as u tends to zero from above.

As with measures of tail correlation, this measure can be extended into higher dimensions. However, there is an issue with this statistic in that it is always zero for the Gaussian copula.

One way round this is to evaluate $C(u,u)/u$ at a finite value of u . This introduces subjectivity, but at least such a measure can be used for all copulas. In the appendix of the paper we have tables for various degrees of freedom of the t copula and the Gaussian copula for various values of correlation and for two, three and four dimensions. The calculation is straightforward for Archimedean copulas where we have a closed mathematical expression.

So far we have been talking about tail association, which is a coefficient, and it is important for issues such as parameterising models. Extreme loss is not a coefficient but a probability. The risk of extreme loss is of more concern for many investors. One key measure, which will probably be familiar to many of you, is the probability of ruin. A related measure, which may not be so familiar, is the economic cost of ruin. The former can be defined as the proportion of observations below the ruin line, that is, the given level of total loss, whilst the latter can be defined as the average value of losses below the line. However, an issue with the risk of extreme loss is how loss is defined.

We have considered two definitions. The first is the loss from current exposure. Here, we consider the maximum acceptable loss from a combination of two, three or even all risks faced by an organisation, defining loss either in absolute terms or as a percentage change.

The second is a little more involved. Here, loss is defined in such a way as to allow the appropriate exposure between groups of risks to be determined. For example, if you are considering a

combination of two risks, you can look at the distribution of losses that you would expect for different mixes of these two risks. In other words, if you are examining a portfolio with only two assets you could look at the probability of ruin for an investment of W_1 in the first asset and $1 - W_1$ in the second, for different values of W_1 . This could also be extended to more asset classes, with the result being an efficient frontier showing how, say, expected profit could be maximised for different probabilities of ruin.

I'd now like to hand back to Mr Sweeting to discuss how this information can be communicated.

Mr Sweeting: In relation to issues such as tail association and the risk of extreme loss, the challenge is not just how statistics should be calculated, but also how large amounts of information should be communicated efficiently. It is also important that the right information is communicated. We believe that this means including not just the strength of any relationship but also its importance. For example, a pension scheme might invest in two property funds whose returns are very closely linked in the tails. This is not necessarily a major risk for the pension scheme if the amount invested in each fund is small. Staying with the pension scheme, it might also be that a significant proportion of the assets is split between an index linked and a conventional gilt fund, which also have a high level of tail association by some measure. However, even though this risk is bigger in terms of the size of the investments, it should not be a concern as both investments act as a hedge for the pension liabilities.

As you can see, this means that there is a lot of information that it is important to get across, and here are some ideas of how to do this.

As I mentioned before, these are just starting points on which other people might like to build.

Figure 8 in the paper is a colour-coded scatter plot. The position of each point on this plot tells us two things. On the vertical axis, it tells us the level of tail association by some measure. In this case it is the coefficient of finite tail dependence. On the horizontal axis, we have the importance of the relationship as measured by the size of the sum of the log sizes of the positions, so we multiply the position sizes and take logs. This means that we should, on the face of it, be concerned with any points falling in the top right hand corner of this chart. These are the combinations of risks that are important in terms of total joint exposure and where there is a high level of tail association.

However, I think we still need a little finessing on this because we need to distinguish between: combinations of matching assets where we do not care if the measure of tail association is high; combinations of assets and liabilities, where we are actually happy if the measure of tail association is high; and all other combinations of risk assets. These are shown in figure 8 as grey, white and black points respectively.

As with the measures of tail association themselves, it is possible to extend this chart to higher dimensions. This simply means that each point represents the combination of three or more risks. However, some thought would be needed on how any colour coding would work or whether it would be better simply to restrict the risk combinations shown to those that really matter.

Another approach is to use a bar chart, such as figure 9 in the paper, rather than a scatter plot. Here, the height of the bars represents the strength of the tail association whilst the line, whose height is measured on the right hand axis, measures the importance of the relationship. The measures and

colour coding are identical to the previous chart. This chart has the advantage that the risks can be ranked either by importance or strength of relationship, but if you are using this sort of approach it becomes a little difficult to interpret if you include all the risks, so you perhaps need to use fewer risks on a single chart.

Next, we have bubble or balloon plots such as in figure 10 in the paper. These will be familiar to many of you as ways of displaying information such as population size or GDP for different countries. We have used it as a way of separating out the position sizes for combinations of risks. So the horizontal axis is the exposure to one risk, for example the value of an investment, and the vertical axis is the exposure to another. The bubbles represent the level of tail association, both in terms of their size and their colour.

Balloon plots can be useful in many circumstances, but they do have some serious drawbacks. Chief amongst them is the fact that the position sizes for many asset classes will be either approximately, or even exactly, the same. This means that many of the bubbles will overlap, making the chart difficult to read.

One way of dealing with this is to instead use a sunflower plot, as in figure 11 in the paper. If all risk sizes are rounded to the nearest 5% of portfolio size, then a sunflower plot can be used to show not only the number of asset combinations for each of these allocations, but also the strength of tail relationships for each one. So looking at figure 11, we have several combinations of risks where one has a weight of 10% of the portfolio and the other has a weight of 50% of the portfolio, the 50% weights being liability values. The number of combinations is shown by the number of petals at this point, and the shade of the petal signifies the strength of the relationship. This means that the position and colour of the sunflowers can be used to give an indication of any risk concentrations.

Those are the main points that we wanted to cover. As I have mentioned, this is not supposed to be a dogmatic view of how we believe this should be communicated but we hope this gives some ideas of how to calculate measures of extreme loss and also how to communicate this information to a range of stakeholders.

Mr E. Varnell, F.I.A. (opening the discussion): First of all, I should like to offer my congratulations to Mr Sweeting and Mr Fotiou on a very comprehensive paper that covers a very wide breadth of material in relation to measures of tail dependence.

I should like to set out some of the things that I took away from the paper which were of particular interest, and offer a few suggestions for improvements in the final draft as well.

Before I do that, maybe I should put my contribution in context. I operate in the insurance sector and currently a lot of us are focused on internal model development and the approval of those models. A lot of that work is connected with understanding the dependency between risk factors and therefore the degree of diversification which can be claimed in reducing capital requirements and demonstrating and ensuring solvency.

However, dependency is not just a subject of interest for model building quants. It is the diversification of risks that underpins the insurance business model. Therefore all insurers and especially their senior management should be interested in the times and situations in which dependency is high and diversification benefit can be lost.

It was, therefore, particularly pleasing to see a paper that not only set out to explain the mathematical complexities of tail dependency, but also sought to explore how these results could be communicated graphically to a less mathematical audience.

So now I want to move onto three things that I took away from the paper. The first of those was the summary of dependence measures. The paper contains a very comprehensive summary of dependence measures. This includes a substantial discussion about empirical dependency measures and some of the issues that can arise in using them. Much of this material is not always found in textbooks

Of particular interest was section 3.2 on concordance, which suggested similarities to the axioms for coherent risk measures that are often used, for example, to justify using T VaR measures rather than VaR measures, as a risk measure.

Perhaps more could have been made of that link, and, in a later version, more explanation could be provided in an executive summary so that the non-technical reader would understand why concordance was worth considering. This would help senior managers to better understand the concept.

It would also be useful to see dependency measures ranked against those concordance axioms so it is slightly clearer why the authors recommended the coefficient of finite tail dependence at the end of section 5.9.

The second thing that I wanted to bring out was the application to asset allocation. Section 6.3 introduced the idea of ruin lines for asset allocation, and also section 7.2 alluded to the use of some visualisations for illustrating the trade-off between risk and return in portfolio analysis.

For me, these sections gave a glimpse of the way that one could link risk management and asset allocation, and so it would be very useful to hear more from the authors in the debate on how they see these being applied in practice.

The third thing I wanted to pick out was visualisations of tail dependencies. I certainly thought that charts 8 to 11 introduced some interesting ideas for how tail dependency can be illustrated, particularly the use of shade, size and shape to increase the information content that we can show in reports showing tail dependence.

For me there were a few places, however, where it was not quite clear that the text reflected the charts being shown. In a subsequent version it might be useful to review this text.

I know that the authors were speaking to other disciplines in the University of Kent about the issue of communicating tail dependencies. So it would be useful to hear their thoughts on whether we as actuaries could pick things up from other disciplines or whether we are really having to deal with a different set of issues.

Next, I should like to cover a few suggested improvements at a high level. Firstly, I think an executive summary would be a very useful addition to this paper, not least as I think there were lots of very important points in the paper which would be very useful for a non-technical audience, and could be drawn out.

The other comment I would make is that the narrative in sections 5.4 to 5.8 could be revisited and expanded with more signposting. For the less technical reader it might be rather hard trying to get through these sections at the moment.

My final comment is that section 6.3 could usefully address the idea of nonlinearity in ruin lines, which is something that we frequently see in the insurance space.

So, in conclusion, I would say that the paper contains a very comprehensive review of tail dependency measures with a good deal of useful content and some interesting ideas on communicating tail dependence. There are also some very promising ideas on asset allocation which would be good to develop.

The paper could benefit from some additional work in clarifying some of the points being made to make them more explicit for the audience and maybe making it a little bit more accessible to senior practitioners from the industry.

So I very much look forward to seeing subsequent versions of the paper and I would like to thank the authors for their efforts to help shed light on this challenging subject.

The Chair: The discussion is now open to the floor. We are looking forward to an interesting and stimulating debate.

Mr A. Hitchcox, F.I.A.: I have now reached the stage in my career where seniority and lack of grey cells means I am not as good at maths as I used to be. So when you say that your measure of tail dependence $C(u,u)/u$ is always zero for Gaussian copulas, is there any way you could expand on what that means? I run a risk portfolio and I would love to know what the result means in the language of risk and reward and profit and loss.

The Chair: I have a question about effective communication. How do you know whether your communications are being received appropriately by the audience?

Mr Sweeting: The first point is to recognise what the tail dependency calculation means. The coefficient of tail dependence is calculated at the limit. So what you are doing is calculating the proportion of observations in the tail divided by the maximum it could possibly be and then reducing that right down to the absolute limit.

When you do that, no matter what the correlation is in the Gaussian (normal) distribution, the coefficient of tail dependence is always zero.

This means that if you have extreme enough observations, which are normally distributed, they are uncorrelated. If the correlation is less than one then at the ultimate of the extreme they are uncorrelated even for a correlation of .999%. That means that the coefficient of tail dependence is useless in trying to discriminate between the tail association between different normal distributions. That is partly because it is a limiting measure. It is looking at something which is quite unrealistic like infinitely large movements.

Mr S. Basak, F.I.A.: In your paper you have used a couple of conditional correlation coefficients. Have you had a chance to look at partial correlation coefficients in which, instead of conditioning

on a part of the data, you condition on another variable and that implicitly gives you dependency like implicit dependency, tail dependency, maybe, between two variables?

Mr Sweeting: It is something which I have looked at but not in this paper. It is something I looked at in the context of asset allocations for, say, pension schemes when they are actually exposed to some other external variable. For example, what is your optimal mix between assets if you are exposed as an entity to some external risk factor?

I suppose this issue falls into the category of developments that other people could make based on the measures that we used. If someone wanted to use, say, our coefficient of finite tail dependence and extend it into that area, then I hope we have given people the basic building blocks to do those calculations themselves. It is difficult to try to cover everything in this paper which is already over 90 pages long.

Mr T. Brenton (a visitor from Metropo): I come from the stockbroking side so I do not look at this in such a technical way, but focus on the communications side. There is a great deal of pressure, an increasing pressure, from regulators to increase clarity over risk. In your capacity as an asset manager do you feel that same pressure from your client base?

I spend a lot of time reading fund prospectuses which basically denote risk as small, medium or large. Is there going to be some sort of movement in the future towards a more complex scoring system of some sort? I realise that it is a long way off at the moment.

Mr Sweeting: Probably not in the retail sector because they have just agreed a new fund classification system which I think does not make any sense to anybody.

In the institutional sector I think it will start to happen. If you look at where some institutions are at the moment, they are already very interested by the concept that maybe asset returns are not necessarily normally distributed. Maybe just a correlation coefficient is not going to give you every measure of risk.

So once we have put across that idea I think we will start looking at the ways that we can actually communicate some of these relationships between the asset classes using these sorts of measures.

Mr H. Sutherland, F.I.A.: I wonder whether this approach, valuable as it is, does not lead us into the trap of having non-technical people, perhaps non-executive directors, or such, saying “Yes, I have seen the analysis. I have had this information communicated to me. I can sign off and say I am aware of the risks that we are facing.” And then the distribution that was used in putting together the model turns out to be inappropriate in some future crisis situation.

Are we at risk of giving people a false sense of security, or at least not communicating the extent to which other scenarios might create problems?

The Chair: That is a fair point.

Mr P. Cronje (student): Following on from that, in the last few years with the global crisis, copulas have received quite a lot of bad press in terms of modelling credit structures and similar situations. A lot of people are quite wary of using copulas just because the word “copula” is linked to global crisis. Is there not a risk of bad publicity in using this kind of approach?

I accept that for technically skilled people this is quite a good and robust approach. But what about the general public? Might they see this as something else that can just go wrong because it has gone wrong before just because it is a copula?

Mr Sweeting: On the risk of being misunderstood, there is always that risk. I think you need to work out what the compromise is between going into enough detail that you start getting close to the true level of risk and showing something which gives you a broad enough idea of what the risks are without giving people a false confidence in the model.

I think there is probably a broader point about whether people rely too much on models and believe that models provide the answer rather than are just used as tools?

Using this approach is no worse than just using correlations. It is giving the same number of statistics but it is concentrating on the most relevant part of the distribution. It still falls into the realms of “Guns do not kill people. People kill people.” You have to rely on the fact that the people using this information are sensible enough to understand the limitations of the model and the statistics.

In terms of copulas and bad press, if that meant that we therefore stop using copulas and said that we will not use a copula, we will use a multi-variant normal model, which is a Gaussian copula with Gaussian margins, I think we need to try to explain exactly what these are. Even the man who developed the model for using Gaussian copulas in CDO pricing at J P Morgan, published something saying, “I am a bit nervous about the fact that everybody seems to be using a Gaussian copula on this type of default model. It was only a proof of concept. You might want to use a more sensible copula in this model.” Again, “Guns do not kill people. People kill people.”

The Chair: Drawing on my experience in general insurance, from a property catastrophe perspective, if you go back to 1992, there was very little industry around capital modelling and how one assessed the risks of property capital re-insurance exposure. Hurricane Andrew gave rise to a lot of capital coming into that marketplace. Then over the last 15 years or so there were huge steps taken to not only improve the data that was available to do modelling but to understand some of the physics behind those events as well. Both helped to educate insurance management teams as to how to take optimal decisions.

I think it is a question of degree as to whether you are betting on the validity of the tail or whether you are using the tail to allow you to assess alternative options.

I certainly think that there is room for a lot more use of this type of communication in insurance businesses more generally. That is not quite the same sector as you were talking about, but I would observe it is not about betting on the validity of the tail, it is more saying, “On this strategy, that looks like 50 times in 10,000 it is going to be better.”

Mr Sweeting: Regarding the extent to which you develop strategies using these sorts of approaches, or use them as a tool, an emphasis on communications aspects is important. If you have a scatter chart you can say, “I am not too happy with all those black dots up in that corner. They are warning flags. They are making me think about what I’m doing” rather than saying, “I’m going to try and optimise my strategy such that I have dots on a particular part of the diagram.”

Mr A. Kaye, F.I.A.: I should like to congratulate the authors on a very comprehensive and interesting paper. With insurance risk data and asset management data one usually has a very large number of data points so one can apply statistical techniques to them with some degree of assurance that they have validity. However, we also have operational risk to worry about, as well as legal risk, fraud, and other types of non-standard risks which can also give rise to extreme loss, as we have seen very recently with UBS.

If we are presenting the types of techniques described in the paper to non-technical senior management and executives, what size of sample space does one have to have, do you think, in order to address the recipients responding, "Okay, you have something very sophisticated here but what does it mean in practice? There does not seem enough data to obtain a real sense of what it means."

How does one impress on the audience that these types of techniques are actually meaningful to them?

The Chair: I think we are probably going to broaden into stress and scenario testing. But let us answer the statistical question first. What sample size do you need?

Mr Sweeting: The true actuarial answer is, "Well, it depends". It depends on the level of confidence you are looking at, so the further into the tail you go, the more observations you need. It is not maybe so much about sample size but about the type of risk. If you are looking at market risks you are likely to have quite a lot of reasonably homogeneous market data and you are probably able to use these sorts of results. If you look at operational risk, sometimes there might be something for which you could use this approach. For example, the losses you might have from small instances of maladministration, where you might have several hundred a year.

If, however, you are looking at UBS-type scenarios, then this approach is going to tell you nothing. For that you should be looking at scenario analysis: you should be looking at the sort of processes used and their weaknesses and concentrating as much on trying to prevent those risks as trying to quantify them. Even then the sort of quantification you will be doing would be to look for the worst case scenario. You would do it that way rather than relying on the sort of graphical and quantitative approach outlined in the paper. This approach is definitely aimed towards risks where there are large volumes of homogeneous data.

The Chair: If there are no other contributions, we will move to asking the authors to respond.

Mr Sweeting: I should like to answer some of the points which Mr Varnell raised earlier which I was not able to do during the discussion. I should like to thank him for his comments on the paper, which are very welcome. Things like extending this to asset allocation work could possibly be a subject for another paper. If you are extending this into a broader area, I see no need to deprive someone else of another fairly lengthy publication.

A lot of the points were well made in terms of the signposting. One of the things that I was quite keen to do with this paper was to come up with solutions which would be understandable and usable by people who were not necessarily deeply technical.

One of the problems we might have is that the paper is quite long so although solutions are in there they might not be that easy to find. So I hope putting them into an executive summary would make it clearer that you do not have to be a rocket scientist to do a lot of the things that we are suggesting in this paper.

In terms of what we found from other disciplines, things like the bubble plot and the sunflower plot were from the statistics department. A lot of what we are looking at here, the relationships between large numbers of different assets, is something which we probably have more of to worry about than many other disciplines. Certainly in the biosciences it is just not something which features highly in the sort of statistics that they need to communicate. So, most of the assistance that we received in that area was from disciplines which were fairly close to our own. The other areas were not able to add quite so much.

Most of the other matters we have answered as we have gone through the discussion. I should like to thank you all for coming.

The Chair: It remains for me to express my thanks to the authors for their contribution to this paper.