Beyond Correlations : The Use and Abuse of Copulas in Economic Capital Calculations

Andrew Chernih[1], Mateusz Maj[2], Steven Vanduffel[34]

Abstract.

Copulas have become a buzzword in recent years in the academic community, and practitioners are paying more and more attention to the choice of a copula in risk management applications.

This paper gives a non-technical and pedagogical introduction to the topic of copulas and explains their role for economic capital calculations.

Risk professionals may be tempted to dress up models by using sophisticated tools like for instance copulas. This is because these toys give them the possibility to give a scientific flavour and "sérieux" to their models, and as such may serve as "an umbrella" towards the different stakeholders involved.

We provide examples to show that models that involve complicated copulas are by no means better than simple but robust and transparent models and do not always add value. However, building a simple as possible, but not too simple, model requires significant actuarial training and expertise.

Keywords: Risk, Dependence, Model

1 Introduction

Copulas have become a buzzword in recent years, especially in the statistical and financial community where they are the topic of numerous papers and discussions.

A copula essentially "joins together" different risks and, as such, is not a rocket science concept, but a rather straightforward statistical concept that allows us to deal with dependence.

Despite extensive literature consisting of many excellent academic and practical papers, there are still many commonly heard fallacies in practice such as:

- "To aggregate risks, all I need are correlations".
- "I assume dependence/independence, hence there is no copula".
- "I need to use a Variance-Covariance approach with tail-end correlations to capture fat tails and tail dependence".
- "The mere existence of the concept of copulas guarantees that economic capital (ECAP) calculations can always be made".

This confusion is partly due to the ubiquitous discussion of "correlations" but infrequent discussion of "copulas" amongst practitioners, and this memo aims at demystifying and clarifying the different concepts.

As we will not discuss what economic capital exactly means, in the remainder of the paper it has to be understood as an upper quantile (e.g. a 99%-quantile) of some random variable representing the random risk at hand.

The remainder of this paper is structured as follows. In Sections 2 till 4 we will discuss correlations, copulas and the relations that exist between these. In Section 5 we discuss the practical use of copulas in ECAP calculations. Finally, section 6 concludes

2 What are linear correlations?

Consider two risks, for example stock returns and interest rates, and assume that we wish to describe the relationship that exists between these two risks. Describing this relationship is about answering questions such as "What happens with the return on the stock market if interest rates increase by 0.1%, 1%,…,5%"; or on the contrary, "what if rates decrease by 3%".

Correlations give some information about this relationship: They can be interpreted as a measure for the average **linear** relationship that exists between pairs of risks.

Average relationship:

The yearly return on the stock market is likely to be unrelated with the mortality risk of a life portfolio, and hence one also expects the average correlation to be zero. Most people will feel entirely comfortable with this.

Let us now compare the yearly stock return with the returns on a portfolio that invests in a combination of a call and a put option on a stock index.

Then, we observe that in case the stock market sharply increases the call option will become more and more "in the money" as well and the value of the put will stay close to zero.
As a result, in this scenario the value of our portfolio will become strongly positively dependent with the value of the underlying stock.

On the other hand, a sharp decrease in the stock market will ensure that the put option will increase in value. Hence sharply decreasing

[1] X-Act Consulting, www.x-act-consulting.com
[2] K.U.Leuven
[3] K.U.Leuven
[4] Corresponding author, Steven.Vanduffel@econ.kuleuven.be

stock markets will go hand in hand with increasing portfolio values showing a negative dependence in this case.

Hence, we may observe that that the (average) correlation is zero although there is a very strong relation between both risks. The reason for this is that the dependence is “switching” between positiveness and negativeness depending on the direction the stock is taking.

Linear relationship:

Most people know what a standard normal risk is, and they will also agree that a standard normal risk X is perfect positively dependent with the risk exp(100*X).

However, it can be easily shown that despite the “perfect dependence” that exists between both risks the correlation is close to zero. The reason for this is that the relation between both risks is not linear but exponential.

“Correlations provide a picture for the average linear dependence between risks but the reality is much more complex and this complexity can, theoretically speaking, be described using copulas”.

The above mentioned features of “linearity” and “averaging” also stress the intrinsic limitations of the correlation concept.

3 What are copulas?

A copula “joins together” different risks. Essentially, they tell us how the risk Y behaves if we know the behaviour of X, and it does so for all realisations of the risk X (‘small’-‘medium’-‘large’).

As such copulas enable one to fully describe the complex relation that exists between different risks.

- An independent copula means that realisations for Y will occur regardless of what is happening with X. Note that an independent copula implies zero correlations but the opposite is false.
- A comonotonic copula means that Y is fully positively dependent with X, i.e. knowing X implies knowing Y with certainty.
- A Gaussian copula - essentially - means that there is a linear relationship between the different risks (after transformation) meaning that they can be written as a linear function of each other plus some random noise also implying that the dependence is equally strong everywhere. In this case we need correlations only.
- Other copulas are the Cauchy copula, Student-t copula, Gumbel copula, amongst others.

In appendix 1 some more information on copulas is provided. An excellent actuarial introduction to the topic can be found in Frees and Valdez (1998).

4 How do copulas link with correlations?

When aggregating risks, one often hears only of “correlations”. This has led many practitioners to believe that only a correlation is required.

In fact, a more correct statement is to say that correlations are required for the aggregation of different risks but are not enough.

Copulas extend the notion of correlations, and it is only when the relationship between the different risks is linear, in the sense as described above, that correlations are enough and then we say that a Gaussian copula may be appropriate.

In appendix 2 we show four different plots of two normal risks X and Y with zero correlations, but using another copula. Before looking at the graphs many people would think that there is no difference between the four situations because in each case we have 2 standard normals with a zero correlation.

Indeed, if we would apply the classical Variance-Covariance formula to the 4 different cases, then we would find that the ECAP, which is essentially an upper quantile, of the portfolio $X+Y$ is the same for each of the different cases.

However after considering the 4 different graphs everybody will agree that, for instance, the third situation (with the Cauchy copula) depicts a large degree of upper tail dependence (if X is large, Y is likely to be large as well) and should give rise to the most conservative economic capital number for the aggregated position $X+Y$.

5 Do copulas have a future in practice?

Economic capital calculations are often extremely sensitive to the choice of copula: Appendix 3 provides an eloquent illustration of this for the case of credit risk ECAP calculations.

Unfortunately, estimating copulas is a difficult exercise, which requires high quality data and that is often not available, or even worse, cannot be available due to the intrinsic low frequency feature of the area at hand (e.g. credit risk).
One may think that for economic capital purposes we only need data that reveals the structure of the upper tail dependence. Whilst it is true that ECAP is most sensitive to upper tail dependence it is precisely in these instances that observations are extremely rare.

In the absence of empirical data fitting a copula becomes a meaningless exercise.

A typical area in which sufficient data is simply not available to even start thinking about the theoretical appropriateness of which copula to use is credit risk, see appendix 3

Also for the aggregation across risk types (life, non-life, credit, market,…) data availability is a significant challenge.

On the other hand, for ALM risk data is readily available. But even then, the estimation of the correlations and the Gaussian copula is not a straightforward exercise at all.

6 Conclusions

We conclude that when the estimation of copulas is a cumbersome exercise (the majority of cases), the corresponding ECAP results have to be considered with extreme caution.

- In these instances we recommend simpler and more transparent but still consistent models for ECAP calculations.
- In this case we also do not recommend using detailed ECAP results to steer involved decision making, such as detailed portfolio optimisation exercises and client profitability analysis.
- Instead, for day-to-day business steering we believe one should focus more on measures for risk that relate to volatility – which is easier to estimate - and less to difficult to measure upper quantiles.
- Discussions regarding the absolute level of ECAP are then relatively meaningless as well.

However, we warn against the use of overly simplistic models: a model must be internally consistent as well.

- For example a (overly) simplistic model for credit risk could consist of the estimation of the portfolio standard deviation (the so called UL). Then for ECAP calculations one may multiply this number with a fixed number. It can be shown that there are portfolios where such an approach can underestimate, with certainty, the ECAP by a factor of at least 10.
- Using a Variance-Covariance framework and so-called tail-end correlations is not a consistent solution either; see appendix 4 for further details.

We believe that the choice of the copula must in the first place be driven by an analysis of the available data and not by the "whistles and bells" of a particular copula. In the second place, especially when data is not really available, the choice should be driven by practical applicability:

- Do we make calculations easier?
- Do we understand what we are doing?
- Do we give the end-users the false idea of being able to provide detailed and accurate numbers when we know that this is not possible?

With regards to economic capital aggregation, the challenge consists of simple but still consistent and well balanced models, and this is not an easy task at all which requires experience and sound training.

ACKNOWLEDGEMENTS

Steven Vanduffel and Mateusz Maj acknowledge the financial support by the Onderzoeksfonds K.U.Leuven (GOA/02: Actuariële, financiële en statistische aspecten van afhankelijkheden in verzekerings- en financiële portefeuilles). Steven Vanduffel also acknowledges the financial support of Fortis (Fortis Chair in Financial and Actuarial Risk Management). The authors are grateful to an anonymous referee for suggestions that helped improve the presentation of the paper.

REFERENCES

[1] Frees, Valdez (1998). Understanding Relationships Using Copulas, North American Actuarial Journal, volume 2, No. 1, 1-25.

[2] Frey, Mc Neil and Nyfeler. (2001). Modelling Dependent Defaults: Asset Correlations Are Not Enough!, Working Paper, ETH.

Appendix 1: Some more information on copulas

Consider two risks X and Y and assume that we are asked to simulate the different possible outcomes of $X+Y$ and the probabilities for each of these events.

When drawing realisations x for the risk X, we first pick a random number u between 0 and 1 and next we transform this into a particular realisation x using the formula $x = F^{-1}(u)$, where F denotes the distribution function of X.

A similar reasoning holds when drawing realisations y for the risk Y. We first pick a random number v between 0 and 1 and next we transform this into a particular realisation y using the formula $y = G^{-1}(v)$, where G now denotes the distribution function of Y.

Next, if we want to know the joint realisations (x,y) it is not enough to put all the realisations $F^{-1}(u)$ and $G^{-1}(v)$ together. Apart from the knowledge about F and G (hence the marginal risks) we also need to know how a drawing u between 0 and 1 for the first risk relates with a drawing v for the second risk. This is effectively described by a copula.

If the drawings from U and V are unrelated then we have an independent copula.

If a drawing from U is also used as a drawing from V, hence u=v, then we have a comonotonic copula which describes the strongest possible dependence.

If the drawings from U and V are related in the same way as what would be needed to generate jointly normally distributed risks then this is a Gaussian copula.

Appendix 2: Link between correlations and copulas

In this appendix we show four situations where the outcomes of two normal risks with zero correlations have been simulated, but another copula is used.

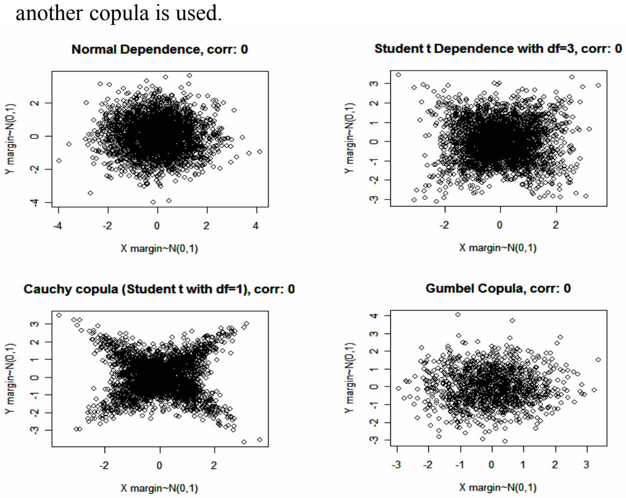

Figure 1: Each graph plots 2 standard normal risks with zero correlation, but each time another copula is used.

Before looking at the graphs many people would think that there is no difference between the four situations because in each case we have 2 standard normals with a zero correlation. Indeed, if we would apply the classical Variance-Covariance formula, then we would find that the ECAP of the portfolio $X+Y$ is the same for each of the different cases.

However after considering the 4 different graphs everybody will agree that, for instance, the third situation (with the Cauchy copula) depicts a large degree of upper tail dependence (if X is large, Y is likely to be large as well) and should give rise to the most conservative economic capital number for the aggregated position $X+Y$

Appendix 3: Ecap results are sensitive to the copula choice

In this section we will illustrate that credit risk economic capital calculations are very sensitive to copula issues.

For Credit Risk calculations financial institutions nowadays have built models to assess the individual risks as well as the correlations that exist between these risks.

Assessing the correlations is akin to measure the likelihoods that pairs of companies will default together. The order of magnitude of these probabilities is typically 1/10,000 or even less.

Then, it is clear that no sufficient empirical default data can ever exist to validate estimates of the correlations that exist between two companies. At best, one is able to assess some portfolio specific correlations.

Copulas go far beyond the notion of correlations. They will address the likelihood that three and more clients will default together but this is virtually impossible to do due to a lack of default data.

For instance, an accurate estimation of the likelihood that IBM, Suez and Deutsche Bank will default is almost like predicting the next winning numbers of the British national lottery. Hence, in a credit context it is not possible to actually measure or validate a copula as there is insufficient default data.

Then, the only way to proceed is to resort to testing all kind of different copulas and this was done by Frey, Mc Neil and Nyfeler (2001), amongst others. They have shown that the model risk related to the choice of the copula in a credit context is enormous.

<u>**Example**</u> *(Frey, Mc Neil and Nyfeler (2001))*

Let us consider a homogenous portfolio consisting of 10,000 bonds, where all default probabilities are identical and where the asset correlation between any two counterparties equals a given constant $\rho > 0$. The other parameters for each bond are as follows:

- Exposure at Default = 1
- Probability of Default = 0.5 %
- Fixed Loss Given Default = 100 %
- Asset correlation ρ=3.80%

In their study the authors test different dependency structures (copulas) whilst preserving the (asset) correlations. They consider the following models for computing a *99% Ecap*:

- Gaussian copula
- t copula with different degrees of freedom *v* (*v=4,10,50*)

Multivariate structure	EL	$Ecap_{99\%}$
Gaussian	50	157
t50	50	261
t10	50	589
t4	50	1074

Table 1: Credit Economic Capital results when varying the copula while preserving the correlations.

From Table 1 we can see that a model that is built using individual credit parameters and the correlations is still subject to significant model risk. The choice of the copula has a massive influence on the risk measures – for example, when we move from a Gaussian model to a t-model with 4 degrees of freedom, our 99% Ecap is inflated by a factor of 7.

Appendix 4: The use of tail-end correlations

The Solvency 2 framework and also many risk professionals appear to continue using a variance-covariance formula, implicitly assuming that risks are all normally distributed and also that there is a Gaussian copula to describe the dependence structure.

It is indeed well-known that a drawback of multivariate normal risks is that they are not designed to capture "fat tails" which are sometimes encountered in practice.

Hence, assuming normal risks may underestimate the tails of the individual risks and also a Gaussian copula is not always adequate to depict the actual (upper) dependence, i.e. the dependence under 'stress scenarios'.

In order to accommodate these two problems people may propose to use a higher value for the correlation coefficient in their Variance-Covariance framework and this correlation is then called a tail-end (or stressed) correlation.

We do not recommend such practice:

- If a 'normal' correlation cannot describe the dependence also a 'stressed' tail-end correlation, whilst giving some feeling of comfort as regards to so-called prudent decision making, cannot describe the dependence.

- Moreover, it is not at all certain that the resulting ECAP calculations are actually conservative. Any variance-covariance framework may, or may not, underestimate the real risk.

- Also, the other risk measures such as the volatility and other quantiles are affected by using these tail-end correlations and their estimation becomes flawed.

- Such an aggregation procedure still assumes normal risks (which are believed to underestimate the risk). We believe this issue should be addressed first, and then one can address the issue of the dependence that exists between the different risks. For example, when aggregating insurance portfolio risk (e.g. lognormal distribution) and credit portfolio risk (e.g. beta distribution) the method that is used should preserve the individual distributions. In contrast, a Variance-Covariance approach imposes normality and hence creates internal inconsistencies.

- If one wants to use a 'stressed' correlation one should also use 'stressed' volatilities and risk premiums, amongst other parameters. In fact, the whole framework should be reconsidered assuming a 'stressed' situation.

- It is not at all clear how 'tail-end correlations' are to be selected. For example, Solvency II QIS3 claims to use 'tail-end correlations' – but then what is the 'average' correlation and how does one go from an 'average' correlation to a 'tail-end correlation'? It is difficult enough to estimate the tails of most distributions with any degree of accuracy – as a result, it would appear near-impossible to be able to specify correlations between tails with any degree of certainty at all.