Tail Risk, Systemic Risk and Copulas

Andy Staudt, FCAS, MAAA

Abstract: Copulas are an elegant mathematical tool for decoupling a joint distribution into the marginal component and the dependence structure component; thus enabling us to model simultaneous events with a greater degree of flexibility. However, as with many statistical techniques, the application of copulas in practice is as much art as it is science. And risk management considerations, such as the increased focus on tail events over central moments, should drive selections of copulas just as much as statistical goodness-of-fit analysis. This paper focuses on several modeling considerations when working with copulas from the perspective of adequately accounting for the behavior in the extreme tails of both the marginal and joint distributions.

Keywords. Copulas; tail risk; systemic risk; joint loss distributions.

1. INTRODUCTION

There is all too often a tendency to focus on what is *reasonably possible* at the expense of what is *remotely probable*. Prospect Theory, pioneered by Daniel Kahneman[1] and Amos Tversky, argues that individuals conflate negligibly or near-zero probabilities with zero probability. When there is a sufficiently remote chance of an event occurring, say 0.01%, most will dismiss this event as even a possibility. However, these remote events are not only likely, but often their likelihood is understated due to a limited understanding of these increasingly small numbers (i.e., what in actuality constitutes a 1-in-10,000 year event when we only have several hundred years of data from which to draw conclusions). And while it may be human nature to ignore such remote probabilities, it is exactly this type of mistake which we, in a risk management context, can not afford to make; as it is these negligible events which can make, or more importantly break, a company. Not only is it essential that we concern ourselves with these unlikely events in isolation, or tail risk, but it is becoming increasingly evident that we also concern ourselves with these unlikely events in tandem, or systemic risk. As the recent financial crisis illustrates, tail and systemic risk are very real and very devastating.

It is now apparent that a major shortcoming in many of the models underlying our financial system is that they failed to adequately comprehend, or just ignored, the risk in extreme events. While it is increasingly in vogue to dismiss many of these models out of pocket, we would argue that it is not the mathematics which are inherently flawed, rather it is the assumptions and simplifications made when implementing such models which are flawed.

[1] Kahneman won the Nobel Prize in Economics in 2002 for his work in this area.

As an illustration, we specifically look at one such model – the copula. The copula is a mathematical tool for modeling the joint distribution of simultaneous events. From the perspective of tail and systemic risk, the copula is interesting in that it allows us to decouple the marginal distribution (that which is associated with tail risk) from the dependence structure (that which is associated with systemic risk) and model each separately with a greater degree of precision. Greater precision, however, does not necessarily ensure greater accuracy. And many copulas, the normal in particular, are unsuitable for modeling extreme behavior. This paper describes several of the considerations in modeling joint behavior with copulas focusing on delineating the choices which will most appropriately reflect the underlying tail and systemic risk – and consequently, the decisions we make.

1.1 Objective & Outline

This paper covers the following areas:

- *Correlation.* Because correlation is easily distorted by outliers and nonlinearities, it may lead to the incorrect calibration of certain copula structures which ultimately impact our measures of risk. Furthermore, because correlation does not provide a roadmap to a unique copula, it may lead to the selection of a copula which does not adequately allow for large losses.
- *Marginal distributions*. Many marginal distributions do not adequately capture the probability of extreme 1-in-*n* year losses and as such understate tail risk. But not only this, misspecifying the marginal distribution may also cause the copula structure to be misspecified, leading to understated systemic risk.
- *Tail dependence*. Tail dependence is a measure of the dependence between two risks in the tail of their joint distribution (i.e., the probability that two companies simultaneously default). To this end, tail dependence can be thought of as a proxy for systemic risk. However, many copula structures do not allow for this type of dependence and as such understate the probability of simultaneous extreme events.
- *(A)symmetry*. While symmetry is common in theoretics, it is rare in nature. However, many of the most popular copulas are symmetric and thus unable to account for the skew associated with many risky, real world events. Asymmetric copulas, on the other hand, do a much better job of modeling these simultaneous extreme events in either or both tails.

Each self-contained section follows roughly the same structure. We first introduce the topic and explain how it relates to tail or systemic risk. We then present an example which uses actual data in topical risk management situations to illustrate the effect certain assumptions have on ultimate measures of risk. Using objective goodness-of-fit criteria, we show that the more conservative models often provide the best fit. Finally, we end each section by offering a general rule of thumb for working with copulas.

1.2 Background & Research Context

Simply put, a copula is a mathematical tool for modeling the dependence structure of a multivariate distribution separate from the marginal distribution without having to explicitly specify a unified, traditional joint distribution. Essentially, copula mathematics are a magnifying glass which allow us to analyze and model with greater precision the dependence relationships between associated random variables. This flexibility means that greater emphasis can be placed on the idiosyncrasies of multivariate distributions, especially with respect to behavior in the extreme tails, leading to models which more accurately account for the entire distribution rather than just the central moments.

In actuarial science, copulas have been used for a variety of purposes including simultaneously modeling loss and allocated loss adjustment expense (ALAE) amounts, measuring the benefit of diversification to multiline insurance products, estimating the default risk of a portfolio of reinsurance receivables, and allocating economic capital by line of business. The following contains practical examples of copulas within a variety of these contexts as a means of illustrating how copula specifications can alter our understanding of risk especially with regards to extreme tail behavior.

There is no shortage of research on copulas, as is true with most mathematics tightly linked with financial markets. The bibliography of this paper is divided into four sections: literature on copulas in actuarial science, survey literature on copulas, computer packages for modeling with copulas, and more esoteric topics with regards to copulas. The purpose of the reference section is to direct the interested reader to literature most relevant for a given purpose.

2. CORRELATION

Correlation, as measured by the Pearson correlation coefficient, has increasingly become a proxy for expressing dependence.[2] In some situations this is appropriate, however, more often, correlation is used in a manner which is inconsistent with its actual meaning. This section explores these situations. First, we detail two problems with correlation as a measure of dependence, namely that (1) it does not necessarily uniquely define the joint distribution and (2) it is distorted by outliers and nonlinearities. Next we present an example which illustrates the consequences of using correlation to specify and calibrate the copula structure.

---

[2] Forthwith, "correlation" refers to Pearson's linear correlation coefficient *rho*.

2.1 The Relationship between the Assumption and the Risk

2.1.1 Correlation does not [necessarily] uniquely define the joint distribution

A classic result in statistics states that independence implies zero correlation, but that zero correlation does not necessarily imply independence (normality is also needed). Even without exploring the meaning of this statement, the logic indicates the problem with correlation—correlation is a weak supposition. Just by virtue of knowing the correlation, we really do not know that much. And thus, it becomes dangerous to assume that by knowing the *correlation*, we truly understand the *dependence* between risks. More specifically, the implication for modeling joint distributions with copulas is that correlation does not necessarily uniquely define the multivariate distribution.

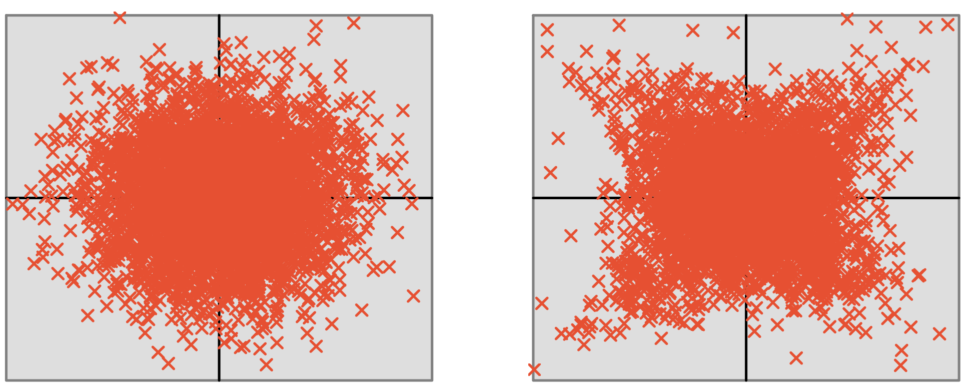

(a) Normal copula (b) *t* copula

Figure 1. Scatterplots of bivariate data generated assuming zero correlation.

Consider Figure 1, based on a similar exposition in Embrechts et al. [20], which compares values simulated from two separate joint distributions. Figure 1(a) shows values which were simulated from a distribution specified by a normal copula and Figure 1(b) shows values which were simulated from a distribution specified by a *t* copula. In both examples, the correlation coefficient is zero. However, this lack of correlation does not necessarily imply that the data is independent. Only the data in Figure 1(a), simulated using a normal copula, is independent. The data in Figure 1(b), simulated using a *t* copula, is in fact dependent. Specifically, there is positive dependence in the tails of the distribution which is not only evident in the grouping of data points in the upper-right and lower-left corners, but can be derived mathematically (and will be in later sections for other purposes). This tail dependence implies that, everything else being equal, the *t* copula might be better suited for modeling joint behavior in situations where systemic risk is of a real concern.

2.1.2 Correlation is easily distorted

Pearson's linear correlation measure is not robust to outliers and because it is a measure of *linear* association it often fails to comprehend the full dependence found in *nonlinear* relationships. As an alternative, rank correlation measures, such as Kendall's *tau* and Spearman's *rho*, are more suitably robust to outliers and, because they operate on ranks of the data rather than nominal values, can both capture nonlinear relationships and are invariant to certain transformations such as the natural logarithm (a very useful technique in modeling probabilities). Furthermore, rank correlations actually have a natural place in copula mathematics, however this is beyond the scope of the paper.

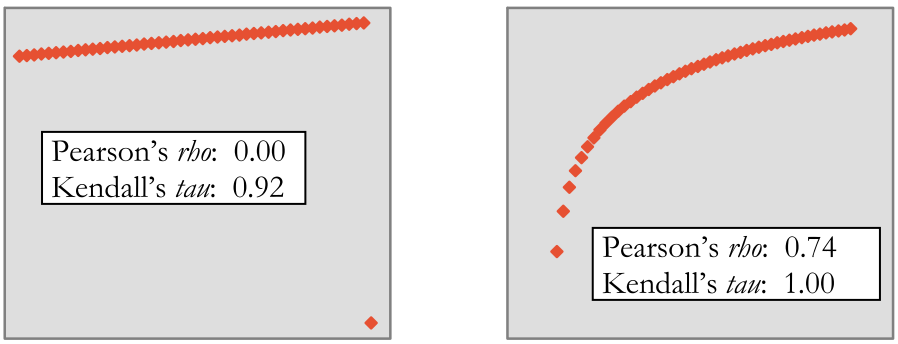

(a) Outliers (b) Nonlinear relationships

Figure 2. Weaknesses in the linear correlation measure.

Figure 2 illustrates these weaknesses. In Figure 2(a), the single outlier in the bottom right completely distorts Pearson's correlation measure while only slightly distorting Kendall's measure of association. In Figure 2(b), the data is generated by an exact, albeit nonlinear, relationship and because Pearson correlation is a measure of linear association it does not recognize the perfect relationship whereas Kendall's *tau* does. In both these examples, Pearson correlation would lead to significantly understated estimates of risk.

Now, *if* we know that the copula describing the joint distribution is either normal or *t*, we can parameterize the copula using either empirical estimates of correlation or empirical estimates of Kendall's *tau*. However, as Figure 2 indicates, it is likely that the empirical estimates of correlation will be distorted by outliers or nonlinearities or both, and thus will not be appropriate to modeling joint relationships.

2.2 An Illustration

The following example uses historical loss ratios for the period from 1986 through 2008 as compiled by the Texas Department of Insurance (TXDOI) for the following lines of business—general liability (GL), commercial automobile liability (CAL), commercial multiple peril-property

(CMP-Property), and commercial multiple peril-liability (CMP-Liability). In order to illustrate the concepts in section 2.1, we calculate the risk for the combined book of business as well as the capital allocation implied by each of several copula structures. Specifically, we compare a normal copula parameterized with correlation, to a *t* copula parameterized with correlation and to a *t* copula parameterized with Kendall's *tau*. In this situation, the first copula does not adequately account for the presence of systemic risk in the data and both the first and second copulas are distorted because the correlation measure is distorted both by outliers and nonlinearities.

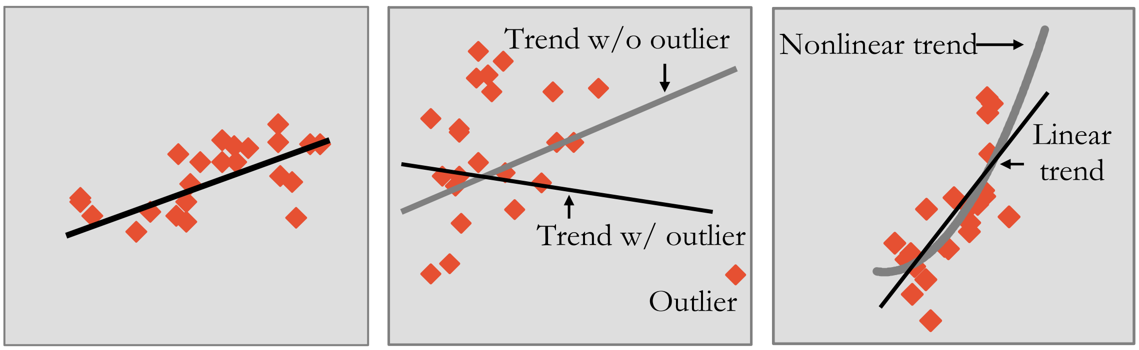

(a) GL by CAL (b) CMP-Property by GL (c) CAL by CMP-Liability

Figure 3. Scatter-plots of loss ratios by line of business.

Figure 3 plots various combinations of these historical loss ratios. The reference trendlines are included to provide a general indication of the correlation (i.e., positive-sloped trendlines have positive correlation, flat trendlines have no correlation and negative-sloped trendlines have negative correlation). From Figure 3(a) it would appear that linear correlation is indeed appropriate for measuring the positive dependence between GL and CAL. However, the correlation statistic between CMP-Property and GL is distorted by the outlier in Figure 3(b); and Figure 3(c) shows evidence of a nonlinear relationship. In both these cases, the correlation measure inadequately expresses the dependence structure, understating the risk.

If we assume that the dependence between these lines of business can be modeled by an elliptical copula, either normal or *t*, then it is possible to quantify the extent to which correlation misspecifies the dependence between lines. Specifically, in the elliptical family of copulas the relationship between Kendall's *tau* and Pearson's correlation *rho* is given as $\varrho=\sin(\pi\tau/2)$. Table 1 compares the correlation implied by Kendall's *tau* statistic using the above relationship with the correlation computed manually.[3] Note that the implied correlation between GL and CAL is identical to the actual correlation as indicated in Figure 3(a). However, the correlation between GL and CMP-

[3] The Kendall Implied Correlation is given as $(2/\pi)\sin^{-1}(\varrho)$ where ϱ refers to Pearson correlation.

Property was indeed distorted by the outlier which produced negative correlation even though Figure 3(b) would indicate a positive dependence without that outlier. The correlation between CAL and CMP-Liability is also understated as the correlation does not recognize the slight nonlinear relationship between these two lines (i.e., the CAL loss ratio increases more sharply for large values of the CMP-Liability loss ratio than for small values).

		Dependence Coefficients			
Line A	Line B	Kendall	Kendall Implied Correlation	Correlation	Cause of Distortion
GL	CAL	0.40	0.60	0.60	Not distorted
GL	CMP-Property	0.15	0.25	(0.10)	Outlier
CAL	CMP-Liability	0.60	0.80	0.70	Nonlinearity

Table 1. Comparison of the correlation implied by Kendall's *tau* statistic in an elliptical family of copulas and the correlation calculated manually using Pearson's *rho* statistic.

In order to further quantify this effect, we fit three copulas to this data. The first copula is a normal calibrated with the empirical correlation, the second copula is a *t* calibrated with the empirical correlation and the third copula is a *t* calibrated with the correlation implied by Kendall's *tau*. Note that comparisons between the first and second copula structures will help to illustrate the premise of section 2.1.1, namely that the correlation matrix does not uniquely define the joint distribution. And that comparisons between the second and third copula structures will help to illustrate the premise of section 2.1.2, namely that correlation is easily distorted.

Table 2 shows the Conditional Tail Expectation (CTE) at the 95th percentile of the excess loss ratio by copula structure for all lines of business combined (for simplicity, we assumed equal exposure by line). Also shown is the percentage capital allocation by line implied by the CTE statistic.

				Capital Allocation				Cramer-von-Mises
#	Copula	Calibration	CTE(95th)	CAL	CMP Liability	CMP Property	GL	Goodness of Fit Statistic*
1	Normal	Pearson's *rho*	1.30	28%	35%	12%	25%	0.11
2	*t* (df=8.5)	Pearson's *rho*	1.35	28%	35%	12%	25%	0.11
3	*t* (df=11.0)	Kendall's *tau*	1.50	28%	40%	10%	22%	0.05

*Smaller values indicate a better fit.

Table 2. CTE at the 95th percentile and percentage capital allocation for each copula structure. The degrees of freedom (df) for the *t* copulas are computed using maximum likelihood estimation holding the copula correlation parameters fixed.

When comparing the first and second copula structures, note that while the percentage capital allocation is not distorted by choosing to use a normal copula, the overall risk is understated. The percentage capital allocation for these copulas is not that affected as these two copulas are parameterized using the same correlation matrix. However, the CTE is higher for the *t* copula as there is greater dependence in the tail of the *t* copula than in the tail of the normal copula.

When comparing the second and third copulas note that both the CTE and the percentage allocation are distorted. Here, we see that the higher implied correlation between CAL and CMP-Liability as well as between GL and CMP-Property significantly drives up the CTE. The percentage capital allocation has also changed as the dependence relationships between lines have changed reflecting the shift in relative riskiness.

Finally, not only is the third copula structure the most conservative, it also objectively, as measured using the Cramer-von-Mises statistic[4], provides the best fit and thus the most reliable estimates of the CTE and capital allocation.

2.3 A Good Rule of Thumb

It is important to remember that correlation is only one measure of association (specifically linear dependence) and as such only tells one side of the story. Although it is useful in defining certain dependence structures (i.e., the multivariate normal distribution), it is easily distorted by outliers and nonlinearities, which can affect the calibration of a copula structure; and it does not provide a roadmap to the correct choice of copula. To these ends, other measures of association, such as Kendall's *tau*, should also be considered as they provide additional insight into the dependence structure and are not as easily tricked by outliers and nonlinearities. Furthermore, other considerations, such as the shape of data, expert opinion and outside estimates of risk, must be weighted carefully and included in any calibration and selection of a copula. Dependence is a dynamic concept, and flat representations like correlation, will always lose something in translation.

3. MARGINAL DISTRIBUTIONS

This section explores the separation between the marginal distributions and the dependence structure in copula models. Although copulas allow us to model these components separately, they are by no means independent of one another. Errors in specifying the marginal distributions can

[4] Generally speaking there are a variety of ways to assess and compare the fits of various copulas. However, most commonly used methods, such as the Cramer-von-Mises statistic, rely on computing some measure of the distance between the estimated copula and the empirical copula.

have far-reaching consequences on the copula fit and the final modeled joint distribution. This is especially true when there is significant tail risk. To illustrate this phenomenon, we fit copulas to historic corn and soybean losses where the marginals are estimated using either a gamma distributions, the empirical distribution function or a mixed empirical-generalized Pareto distribution. We then show how the gamma distribution and the empirical distribution function lead to copula parameterizations which understate the systemic risk relative to the mixed generalized Pareto distribution.

3.1 The Relationship between the Assumption and the Risk

Perhaps the major benefit of copulas is that the dependence structure (i.e., the copula) can be decoupled from the marginal structure and modeled separately. For example, rather than approximating the joint distribution of two risks with a multivariate normal, we can use the copula framework to instill a more precise structure by specifying gamma and lognormal marginals coupled with a Gumbel copula (as shown in Figure 4).

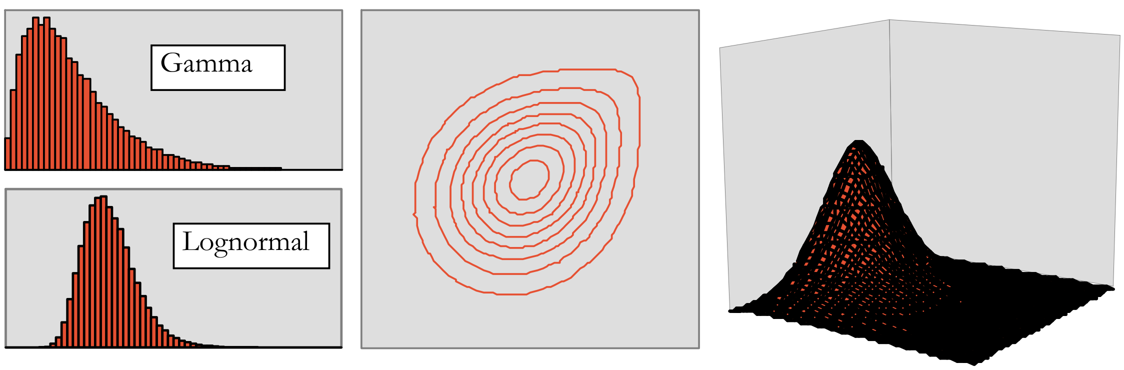

(a) Marginals (b) Gumbel copula (c) Joint distribution

Figure 4. Decomposition of joint distribution into marginal structure and dependence structure.

However, just because the joint distribution can be decomposed into these component parts, does not mean that these component parts are independent of one another. In fact, they are very much linked especially when fitting a copula to data.

There are a variety of ways to fit copulas to data. One of the more popular methods, Inference Functions for Margins (IFM), is a two-step procedure whereby first distributions are fit to the marginals and then maximum likelihood is used to estimate the copula parameters conditional on the marginals fit in the first step. Because this process is order dependent, any misspecification of the marginals in step one will distort the fit of the copula in step two and ultimately the joint distribution. From a risk management context, we should be most wary of marginal distributions which do not appropriately allow for the possibility of extreme events (i.e., tail risk); from a statistical context, we should be further wary of misspecifying these marginal distributions as the error compounds causing us to *also* often underestimate the likelihood of extreme simultaneous events (i.e., systemic risk).

The next section illustrates this "ripple-effect" by comparing both the fit and degree of risk associated with copulas parameterized using various underlying marginal distributions.

3.2 An Illustration

The following example uses data compiled by the Risk Management Agency (RMA) of the United States Department of Agriculture (USDA) for the Federal Crop Insurance Corporation (FCIC). Specifically, we looked at historical corn and soybeans losses (relative to net insured acres) in monthly increments for the period from 1989 through 2008. We show how incorrectly specifying the marginal distribution leads to errors in the calibration of the copula function and ultimately results in CTEs which are understated and an overstated benefit to diversification[5].

This dataset is interesting from a number of perspectives. There is positive dependence between corn and soybean losses (due to common causation by perils such as excess moisture or drought). There is systemic risk (i.e., a peril which completely wipes out a soybean crop in a certain location is also very likely to completely wipe out a corn crop). And there is also evidence of tail risk in the humps, or fat right tails, of the kernel densities fit to historic corn and soybean losses (see Figure 5). It is the fat tails of these marginal distributions which we are most interested in modeling for the time being.

---

[5] Here, the benefit to diversification is specifically defined as the difference between the sum of the conditional tail expectations and the conditional tail expectation of the sum. This statistic measures the benefit to diversifying with lines of business that are not perfectly correlated. Note that the former statistic does not allow corn (soybean) losses in excess of expectations to cancel with soybean (corn) losses less than expectations. And vice versa. However, the later statistic does and thus the difference provides one measure of diversifying with lines of business that are not perfectly correlated.

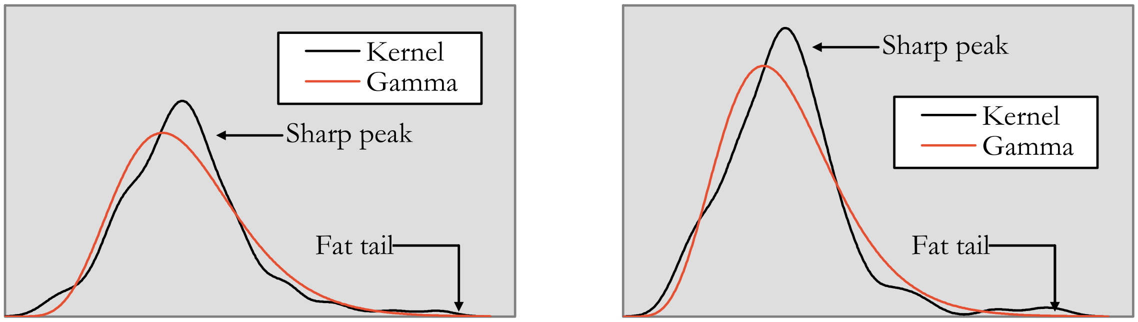

(a) Corn (b) Soybeans

Figure 5. Comparison of kernel density with gamma density fit to historical corn and soybean losses.

Figure 5 also plots the parametric gamma distribution. Note that this distribution does a poor job mimicking the shape of the data. In order to fit the heavy tail, the two-parameter gamma distribution is forced to contort its shape and in the end equally underfits both the sharp mode and the extreme right tail. Further note that the empirical distribution function may fit the data too closely, degrading its predictive power. However, more importantly, although it is not that evident from these graphs, the empirical distribution function does not adequately assign probabilities to values in excess of the maximum observation in the sample (and kernel density estimates often do a poor job of extrapolation). To this end, the empirical distribution function may not be suitable for modeling tail risk.

To address this later consideration, extreme value distribution, such as the GPD, are often mixed together with another more traditional probability distribution and used to model events in excess of a certain threshold (usually set at a large quantile such as the 90th or 95th). This allows us to account for large losses which may not have been occurred historically but are still expected to be a real possibility in the future.

Figure 6 plots the pseudo-observations of the cumulative probabilities based on either a gamma, empirical, or mixed empirical-GPD fit.[6] The lower panels magnify the area in excess of the 90th percentile (i.e., the observations in the joint right tail). Note that the largest observations in the data, when mapped using the selected gamma distribution, are assigned cumulative probabilities very near to one. For the empirical distribution function, the cumulative probabilities are pushed away from one toward the left corner and for the mixed empirical-GPD these cumulative probabilities are pushed substantially away from one. Essentially, the empirical distribution and the mixed empirical-

[6] Pseudo-observations are the actual observations mapped onto [0,1] using the selected cumulative density function of the marginals.

GPD are assigning greater *survival* probabilities to the observed data in the tail.[7] These larger survival probabilities imply the possibility of observations much larger than that seen in the sample and thus allow for increased tail risk.

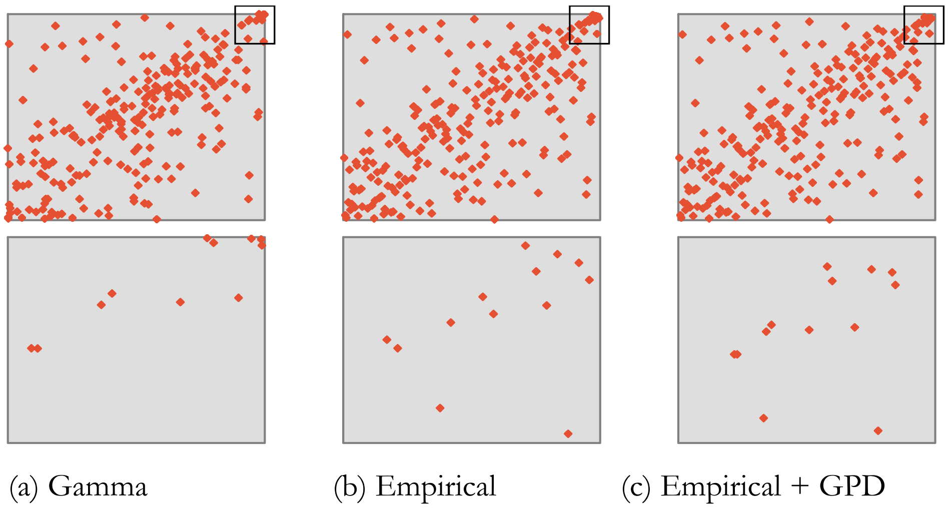

Figure 6. Top panels show pseudo-observations of corn (x-axis) by soybean (y-axis) losses; bottom panels show magnification of 90th percentile excess.

However, these marginal distributions not only affect our estimates of tail risk, but they also affect the calibration of the copula and ultimately our estimate of systemic risk. Table 3 highlights these results. In this example, the copula parameter is larger for the empirical distribution than the gamma distribution and it is larger for the mixed empirical-GPD than the empirical distribution. This implies increased dependence especially with regard to dependence in the tail of the joint distribution. The CTE is significantly larger for both the empirical distribution and the mixed empirical-GPD distribution as this reflects not only the increased systemic risk, but also the increased tail risk in the marginals. The benefit to diversification is also overstated for the gamma marginals as compared to the mixed empirical-GPD marginal. This is because the possibility of simultaneous tail events greatly reduces the actual benefit from diversifying across these random events.

[7] The survival probability is the probability that a random variate takes a value in excess of a given threshold.

Marginals	Copula	Copula Parameter	CTE(95^{th})	Benefit to Diversification	Cramer-von-Mises Goodness of Fit Statistic*
Gamma	Gumbel	1.88	58.7	5.7%	0.036
Empirical	Gumbel	1.89	82.4	5.6%	0.035
Mixed Empirical-GPD	Gumbel	1.93	106.6	4.8%	0.031

**Smaller values indicate a better fit.*

Table 3. Comparison of copulas fit using the inference functions for marginal approach and various marginal distributions.

Again note that the copula based on more conservative estimates of the underlying marginals (i.e., the mixed empirical-GPD) provides the best fit and thus the more accurate estimates of the actual CTE and benefit to diversification.

3.3 A Good Rule of Thumb

For a variety of reasons, including the rigidity of many parametric distributions as well as the poor job historical data does at capturing the future potential of extreme events, many marginal distributions do not allow for a sufficiently high possibility of large 1-in-*n* year type losses. However, not only do these distributions fail to adequately capture the tail risk, but they also distort the calibration of the copula structure in effect understating the systemic risk. To this end, in order to correctly allow for both tail and systemic risk, it is often advisable to use, or at least consider, an extreme value distribution to model losses above a certain threshold while modeling losses below that threshold with a traditional probability distribution.

4. TAIL DEPENDENCE

This section explores the concept of *tail dependence*. Tail dependence is a specific, asymptotic measure of the dependence between two random variates in the tail of their joint distribution.[8] However, it can be more generally thought of as a good proxy for systemic risk. What is most interesting about the tail dependence statistic is that the normal copula, for all nontrivial

[8] Specifically, tail dependence alludes to the probability that a random variable Y takes a value in the extreme tail of its distribution given that another random variable X has also taken a value in its extreme tail (i.e., consider the scenario where X and Y measure bankruptcy for two companies and both companies simultaneously go bankrupt). Mathematically, the following describes the joint upper tail dependence of random variates X and Y: $\lim_{\alpha \to 1} P(Y > F_Y^{-1}(\alpha) \mid X > F_X^{-1}(\alpha))$.

parameterizations, has no tail dependence. This section explores how systemic risk can be understated when using a normal copula rather than some other copula structure.

4.1 The Relationship between the Assumption and the Risk

Dependence, as discussed in the section 2, is a measure of association between two or more random variables over their entire range. Tail risk, as discussed in section 3, refers to the likelihood and amount of loss in the extreme tails. Tail dependence or systemic risk, however, more pointedly measures the *association* in the *extreme tails* of the *joint* distribution. In this regard, tail dependence is not the same as dependence. It is possible for two random variables to be dependent, but for there to be no dependence in the tail of the distributions. This is exactly the situation described by the normal copula and referenced in Figure 1.

Figure 7 illustrates this concept by plotting bivariate random observations generated from copulas fit to daily stock returns for two large reinsurers over the period 1996 through 2008. The graphs have been divided into quadrants where the lower left quadrant represents simultaneous extreme, downward stock movements (i.e., systemic risk). Even though the copulas were fit to the same data, the normal copula produced no joint extreme events. On the other hand, as can be seen from the graphs below, the *t* copula produced three and the Clayton copula produced about five. Note also that the density of the plotted points for the normal copula, as compared to the *t* or Clayton copulas, thins out considerably and quickly as the simulated observations tend toward the left.

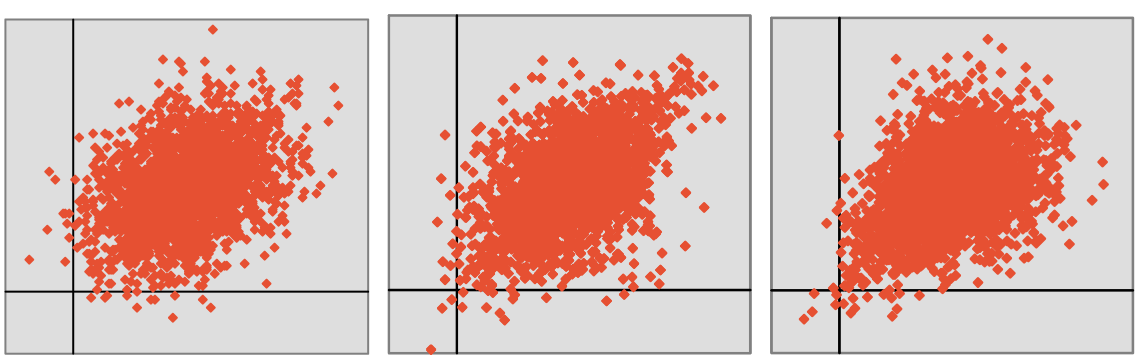

(a) Normal Copula (b) *t* Copula (c) Clayton Copula

Figure 7. Plots of random observations generated from various copulas.

Table 4 shows the upper and lower tail dependence statistics for these copulas. Even though these returns show association as measured by Kendall's *tau*, for the normal copula the tail dependence is zero (whereas it is positive for both the *t* and Clayton copulas). Furthermore, note that the Clayton copula puts the entire tail dependence in the lower left tail whereas tail dependence is symmetric for the *t* copula.

Copula	Kendall's *tau*	Tail Dependence Lower	Tail Dependence Upper
Normal	0.25	0.00	0.00
t (df=4.45)	0.25	0.17	0.17
Clayton	0.25	0.35	0.00

Table 4. Tail dependence statistics for various copulas.

4.2 An Illustration

Suppose we are interested in estimating the default risk of a portfolio of two, million dollar reinsurance recoverables. Assume there is 100% loss on default (i.e., Bernoulli marginals) and that the probability of default is approximately 3.0%. Also assume that the dependence can be described using a Kendall statistic of 0.25. We fit four different copula structures to this data—the normal copula as well as three members of the extreme value family of copulas which all have strong upper tail dependence (approximately 30% in this situation).

Table 5 compares the probability distribution of defaults across the various copulas. Here, the probability that both reinsurers simultaneously default is about 2.5 times as large with the extreme value copulas than with the normal copula. This is because the extreme value copulas allow for a greater possibility of joint default (i.e., simultaneous extreme events). Further, the probability that neither company defaults is also higher with the extreme value copulas. The immediate implication is that one-parameter copulas, of which all of these are, may not be versatile enough to capture the more complex relationships between jointly distributed random variates. In this specific hypothetical, no copula can be said to be "most correct," instead it is necessary to assess not only the input parameters (i.e., Kendall statistic of 0.25), but the output probabilities (i.e., 94.4/5.2/0.6 vs. 95.0/4.0/1.0) as well for reasonableness.

		Extreme Value Copulas		
Probability of:	Normal Copula	Galambos	Gumbel	Husler Reiss
No Defaults	94.4%	95.0%	95.0%	95.0%
One Default	5.2%	4.0%	4.0%	4.0%
Both Default	0.4%	1.0%	1.0%	1.0%

Table 5. Probability distribution of defaults.

Table 6 compares the CTE at various thresholds. While the CTE is approximately equivalent at the lower thresholds, it grows increasingly fast for the extreme value copulas. This is because the extreme value copulas model a higher percentage of joint defaults than would be the case with the normal copula. Without commenting on the appropriateness of one copula over another, the extreme value copulas allow us to be more conservative when estimating the possibility of default, which might just be a good thing.

		Extreme Value Copulas		
Threshold	Normal Copula	Galambos	Gumbel	Husler Reiss
50^{th}	120K	120K	120K	120K
75^{th}	240K	240K	240K	240K
90^{th}	600K	600K	600K	600K
95^{th}	1.10M	1.20M	1.20M	1.20M
97.5^{th}	1.16M	1.39M	1.40M	1.40M
99.9^{th}	1.41M	1.97M	1.98M	1.97M

Table 6. CTE at various thresholds.

4.3 A Good Rule of Thumb

There is too often a tendency to focus on the central moments and distribution of data while ignoring behavior in the tails. In a risk management context, this tail behavior is often the most important driver of results and as such should be given a great deal of care. Where it is possible to get good estimates of tail dependence coefficients, these should be included in the selection and calibration of copulas. If this is not possible, due consideration should be given to the nature of the data specifically with regards to expected behavior in the tails of the distribution. This expert opinion should then serve as much of the basis for the final copula structure.

5. (A)SYMMETRY

Copulas are either symmetric or not – this section focuses on the relationship of symmetry with tail risk (associated with univariate asymmetry) and systemic risk (associated with multivariate asymmetry). Two of the most commonly used copulas, the normal and the *t*, are both symmetric and as such behave identically in the left tail as in the right tail. However, in a risk management context, it may not be ideal to model extreme negative outcomes in the same manner as with extreme positive outcomes. More often, positive outcomes may be associated with general run-of-the-mill probabilities whereas negative outcomes are associated with the unlikely 1-in-*n* year events. Modeling these opposite tails in a similar manner will generally lead to undervaluation of the true risk as both the tail risk and systemic risk will generally be understated.

The concept of kurtosis is also discussed within the context that distributions and copulas which are leptokurtic (i.e., have a higher peak of probability around the mean as well as fatter tails) are more risky.

5.1 The Relationship between the Assumption and the Risk

The normal distribution is an extremely elegant formulation which, because of its mathematical properties, appears again and again in theoretical statistical research. However, it appears less in the real world, as most empirical data just doesn't behave that nicely. Most statistical tests of normality (i.e., does the data follow a normal distribution) are based on the skewness statistic and the kurtosis statistic. Skewness measures symmetry about the mean with the normal distribution being symmetric. Kurtosis, or more specifically excess kurtosis, measures the peakedness of a distribution relative to the normal distribution. Excess kurtosis statistics greater than zero imply a sharper peak in probability around the mean as well as fatter tails than the normal distribution (i.e., increased tail risk).

With that said, two of the most popular copulas are still both symmetric – the normal and the *t*. Perhaps the major criticism of the normal distribution is that there is no tail dependence and thus it is not appropriate for modeling extreme events. However, because it is symmetric it is often not appropriate for modeling most real-world events, many of which tend to have an unlimited downside with only a limited upside. Furthermore, the tails of the normal distribution are considered to be rather thin (i.e., there is a low probability of events at large distances away from the mean). On the other hand, while the *t* copula has both fatter tails than the normal distribution (i.e., positive excess kurtosis) and nonzero tail dependence, it is still symmetric about the mean. In fact, perhaps the major criticism of the *t* copula is that there is only one parameter, specifically the degrees of freedom, which can be used to model the tail dependence.

Figure 8 shows the symmetry of the normal and *t* (as well as Frank) copulas by plotting the probability contours (i.e., 2-D representations of the 3-D probability similar to that shown in Figure 4(c)).

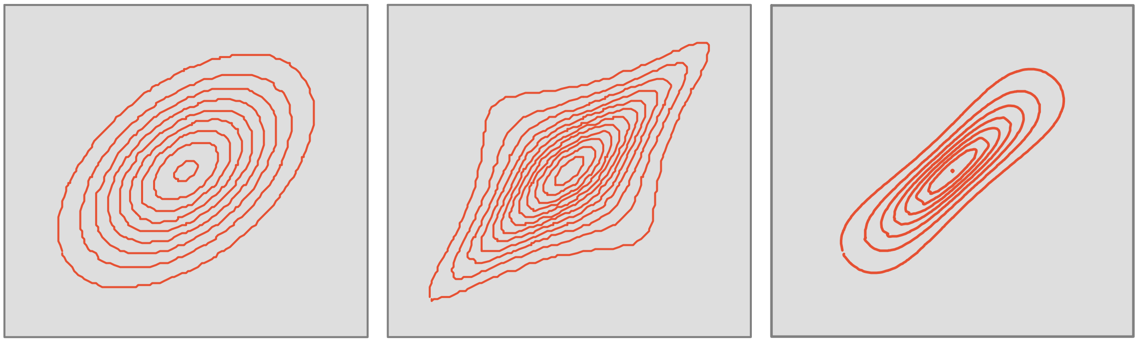

(a) Normal copula (b) *t* copula (c) Frank copula

Figure 8. Common symmetric copulas.

Figure 9 shows contour plots of some common asymmetric copulas. The first two, the Galambos and the Husler-Reiss, are both members of the extreme value family of copulas (along with the Gumbel copula referenced elsewhere) which are characterized both by strong upper tail dependence and right skew. Conversely, the Clayton copula, shown in Figure 9(c), is also an asymmetric copula, but it is instead left-skewed with strong lower tail dependence (and zero upper tail dependence).

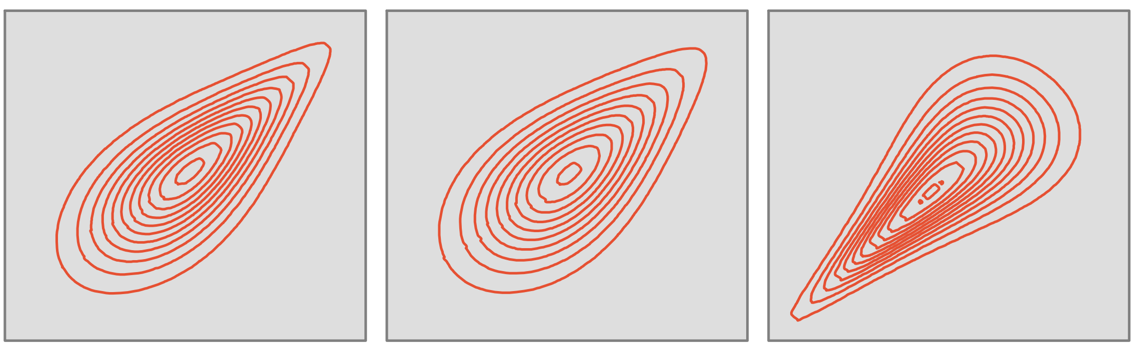

(a) Galambos copula (b) Husler-Reiss copula (c) Clayton copula

Figure 9. Common asymmetric copulas.

5.2 An Illustration

The following example uses data compiled by the Florida Office of Insurance Regulation (FLOIR) on the loss and allocated loss adjustment expense (ALAE) associated with medical professional liability (MPL) closed claims for the period from 2000 through 2009. We show how symmetric copulas do a poor job of fitting the empirical skewness and kurtosis of the data and thus understate the risk.

Figure 10 plots the log of loss amounts by the log of ALAE amounts. There is a definite positive dependence between loss and ALAE amounts (i.e., as loss amounts increase, so generally do ALAE amounts), however what is more interesting is the presence of both a strong right skew and upper tail dependence (contrasted with weak lower tail dependence). Although it is not entirely evident from the graph, this data is also extremely peaked, or leptokurtic, meaning that this data has fatter tails than the normal distribution (i.e., increased tail risk).

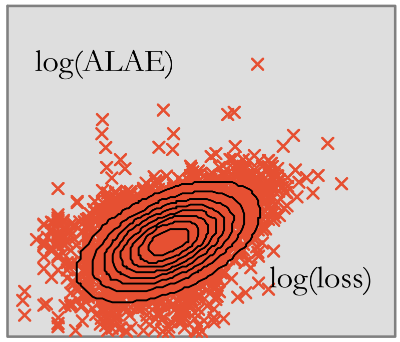

Figure 10. Scatterplot of the natural logarithm of loss and ALAE amounts. A contour plot of the normal copula fit to this data has been overlaid to show how the normal copula fails to adequately capture the shape of the data (i.e., strong right skew and loose left dependence).

Figure 10 makes sense given the possible nature of the data. There is a looser relationship in the lower left corner (i.e., weak lower tail dependence) as small loss payments may be associated with either a constant ALAE per small loss or a large amount of expense perhaps associated with defense which then resulted in a small payment. We would also expect there to be large variations in loss amounts given small ALAE amounts as many claims settle relatively painlessly regardless of the size of loss. There is a tighter relationship in the upper left corner (i.e., strong upper tail dependence) as very large loss amounts are generally associated with very large ALAE payments. Further, note the increased density of points around the median loss and ALAE amounts which gradually taper off in the direction of the upper right corner. This is consistent with a positive excess kurtosis and right skew, respectively, both of which imply increased systemic risk.

To measure the skewness and excess kurtosis, we use Mardia's multivariate extensions of the common skewness and kurtosis statistics. The normal copula will generate values of zero and zero. Table 7 compares the actual skewness and kurtosis of the data against various copulas fit to the data. Note that the skewness statistics for the symmetric copulas are zero and for the asymmetric copulas are greater than zero indicating a right skew. None of the traditional copulas provide a particularly good fit to the data with respect to capturing the risk in the right tail. In all of these examples, the copulas are unable to fully capture the extreme multivariate behavior of the underlying data and as such will understate the ultimate risk. The skew *t* copula does slightly better but still understates the skewness while overstating the kurtosis.

Copula	Symmetry	Skewness	Excess Kurtosis
Actual	**Asymmetric**	**0.50**	**1.50**
Normal	Symmetric	0.00	0.00
Frank	Symmetric	0.00	0.10
t	Symmetric	0.00	0.25
Galambos	Asymmetric	0.10	0.15
Gumbel	Asymmetric	0.10	0.25
Skew *t*	Asymmetric	0.40	1.80

Table 7. Multivariate skewness and excess kurtosis statistics of copulas fit to the log of loss and ALAE amounts. The actual skewness and kurtosis are included as reference.

5.3 A Good Rule of Thumb

To some extent, many of the concepts we rely upon when modeling univariate distributions apply just as well when modeling multivariate distributions. Specifically, if we do not use the normal distribution to model loss severities because the normal distribution is not skewed, why should we use it to model multivariate loss severities. Put another way, because multivariate structures are more difficult to conceptualize than univariate structures, it may often be easiest to think about multivariate modeling in terms of univariate best practices.

Further, it is important to note that although we may often rely on data to determine the ultimate shape of our curves, with copulas the ultimate shape is perhaps more a product of theoretical considerations than it is of data parameterization. As such, prior to fitting copulas to data, it is necessary to take a step back and decide on which copula(s) – symmetric or not, kurtic or not – have a natural interpretation and make sense given any prior knowledge of the risks.

6. CONCLUSION

Perhaps even more so than other statistical techniques, the application of copulas is often more art than science. There will generally never be that one obviously correct answer; however, there are often many wrong answers. More specifically, there are many copula structures which fail to adequately account for the behavior in the extreme tails of univariate and multivariate loss distributions and as such greatly understate the tail and systemic risk. This paper has highlighted several considerations with regard to more appropriately capturing both the tail and systemic risk, including using measures of association more robust than linear correlation, using extreme value theory to model the marginals, selecting a copula which appropriately captures the tail dependence and accounting for the skewness and kurtosis of the underlying data.

However, perhaps most important in the selection of the copula, and the best rule of thumb, is to select a copula which has a natural interpretation (i.e., it makes sense and can be explained) and is consistent with expectations of risk remembering always that the future will never be quite as simple as the past.

Acknowledgment

I would like to thank Rasa McKean, Phil Kane and Ed Yao for their thorough review. This paper is much stronger because of their helpful suggestions and commentary.

7. REFERENCES

Copulas in Actuarial Science

[1] J. Barnett, R. Kreps, J. Major, and G. Venter, "Multivariate Copulas for Financial Modeling," *Variance*, **2007**, Vol. 1, No. 1, 103-119.

[2] S. Desmedt, and J. F. Wahlin, "On the Subadditivity of Tail Value at Risk: An Investigation with Copulas," *Variance*, **2008**, Vol. 2, No. 2, 231-252.

[3] E. Frees, "Understanding Relationships using Copulas," *North American Actuarial Journal*, **1998**, Vol. 2, No. 1, 1-25.

[4] G. Venter, "Tails of Copulas," *PCAS*, **2002**, Vol. LXXXIX, 68-113.

[5] G. Venter, "Quantifying Correlated Reinsurance Exposures with Copulas," CAS *Forum*, Spring **2003**, 215-229.

[6] S. Wang, "Aggregate of Correlated Risk Portfolios: Models and Algorithms," *PCAS* **1998**, Vol. LXXXV, 848-939.

Surveys of Copula

[7] E. Bouyé, V. Durrleman, A. Nikeghbali, G. Riboulet, and T. Roncalli, "Copulas for Finance: A Reading Guide and Some Applications," working paper, **2000 ,** (March 7, 2000). Available at SSRN: http://ssrn.com/abstract=1032533.

[8]

[9] P. Embrechts, R. Frey, and A. McNeil, *Quantitative Risk Management: Concepts, Techniques and Tools*. **2005**. Princeton University Press, Princeton.

[10] P. Trivedi, and D. Zimmer, "Copula Modeling: An Introduction for Practitioners," *Foundations and Trends in Econometrics*, **2005**, Vol. 1, No. 1, 1-111.

Computing with Copulas

[11] R Development Core Team (2009). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, URL http://www.R-project.org.

[12] A. Azzalini, "SN: The Skew-normal and Skew-*t* Distributions," *R Package*, **2010**, Version 0.4-15, URL http://cran.R-Project.org/web/packages/sn.

[13] I. Kojadinovic and J. Yan, "Copula: Multivariate Dependence with Copulas," *R Package,* **2010**, Version 0.9-3, URL http://cran.R-Project.org/web/packages/copula.

[14] A. McNeil and S. Ulman, "QRMLib: Provides R-language Code to Examine Quantitative Risk Management Concepts," *R Package*, **2010**, Version 1.4.5, URL http://cran.R-Project.org/web/packages/QRMlib.

[15] D. Wuertz, et al., "fCopulae: Rmetrics – Dependence Structures with Copulas," *R Package*, **2009**, Version 2110.78, URL http://cran.R-Project.org/web/packages/fCopulae.

[16] D. Wuertz, and Y. Chalabi, "timeSeries: Rmetrics – Financial Time Series Objects," *R Package*, **2010**, Version 2110.87, URL http://cran.R-Project.org/web/packages/timeSeries.

[17] J. Yan, "Enjoy the Joy of Copulas: With a Package copula," *Journal of Statistical Software*, **2007**, Vol. 21, No. 4, 1-21.

Specific Copula Topics

[18] M. Bargest, H. Cossette, and E. Marceau, "TVaR-based Capital Allocation with Copulas," *Insurance: Mathematics and Economics*, **2009**, Vol. 45, No. 3, 348-361.

[19] S. Demarta, and A. McNeil, "The *t* Copula and Related Copulas," *International Statistical Review*, **2005**, Vol. 73, 111-129.

[20] P. Embrechts, A. McNeil, and D. Straumann, "Correlation and Dependence in Risk Management: Properties and Pitfalls," in *Risk Management: Value at Risk and Beyond*, ed. M. Dempster, Cambridge University Press, Cambridge, 176-223.

[21] D. Li, "On Default Correlation: A Copula Function Approach," *Journal of Fixed Income*, **2001**, Vol. 9, 43-54.

[22] K. Mardia, "Measure of Multivariate Skewness and Kurtosis with Applications," *Biometrika*, **1970**, Vol. 57, No. 3, 519-530.

[23] T. Mikosch, "Copulas: Tales and Facts," *Extremes*, **2006**, Vol. 9, 3-20.

[24] R. Shaw, "The Modelling of Reinsurance Credit Risk," *GIRO*, **2007**.

Abbreviations and notations

ALAE, allocated loss adjustment expense
CAL, commercial automobile liability
CMP-Property, commercial multiple peril (property portion)
CMP-Liability, commercial multiple peril (liability portion)
CTE, conditional tail expectation
df, degrees of freedom
FCIC, Federal Crop Insurance Corporation
FLOIR, Florida Office of Insurance Regulation
GL, General Liability
GPD, generalized Pareto distribution
IFM, inference functions for margins
MPL, medical professional liability
RMA, Risk Management Agency
TXDOI, Texas Department of Insurance
USDA, United States Department of Agriculture

Biography of the Author

Andy Staudt is a consultant with Towers Watson in London. He is a Fellow of the Casualty Actuarial Society (CAS), a member of the American Academy of Actuaries (AAA) and has a Masters in Statistics from the University of California at Berkeley. He can be reached at andy.staudt@towerswatson.com or (+44) 020 7170 3479.