

The ART of dependence modelling: the latest advances in correlation analysis.

Peter Blum*, Alexandra Dias†, Paul Embrechts

Department of Mathematics
ETH – Swiss Federal Institute of Technology
CH–8092 Zürich
Switzerland

Abstract

Both at the design stage as well as at the pricing stage of Alternative Risk Transfer (ART) products, the notion of low (zero) beta plays an important role. By now it is well known that for these non-standard products, the interpretation of dependence through linear correlation (and hence the portfolio-beta language) becomes dubious. We review some of the new tools (like copulas) to handle the measurement of dependence in ART products. An example will be discussed.

1 Introduction

From a methodological point of view, Alternative Risk Transfer (ART) for instance aims at securitising risk from a typically more specialised market (the property insurance market, say) to the broader market of finance. Several contributions of the present volume enter much more into detail on the precise definitions used, the marketing, the actual engineering and the pricing and hedging of these products. Our contribution mainly concentrates on some of the methodological issues underlying the pricing and portfolio theory of specific ART transactions. Implicitly in many of our statements, the problem of model risk is present, i.e. the risk encountered in ART pricing and hedging when using “wrong” models. We have hyphenated “wrong” as the whole field of model fitting is a diffuse one indeed. An excellent example

*Supported by Converium Reinsurance Ltd, Zürich, Switzerland.

†Supported by Fundação para a Ciência e a Tecnologia and Faculdade de Ciências e Tecnologia, Universidade Nova de Lisboa, Portugal.

of a model risk analysis for a specific ART product is to be found in Schmock (1999) and Gisler and Frost (1999). For a more general discussion on model risk within finance, see Gibson(2000).

From a marketing point of view, ART products are sold to the portfolio manager using arguments like:

- they are products from markets very different from the usual financial markets;
- they have low correlation with (even termed “independent from” by some of the sellers) other traditional instruments leading to low beta or even zero beta from a portfolio theory point of view;
- they are to be included in a portfolio in order to achieve a better diversification effect, and
- finally, they have lower risk and yet yield higher returns: they hence pierce the traditional efficient frontier.

Products typically falling under such a sales pitch for instance include funds-of-funds (also referred to as hedge funds or non-traditional funds), private equity, securitised insurance risk, weather and energy derivatives. To these, and others, some of the above diversification arguments may apply over some periods of time. We have however also learned that the diversification effect may break down especially in periods of financial distress; the key example coming from the hedge fund crisis in 1998 triggered by the collapse of LTCM.

However one looks at the history of ART, at the basis of both its successes as well as its failures lies the notion of dependence; when the returns on an ART product are (close to) independent from the returns of other portfolio instruments, it pays to add these products to one’s investment portfolio. We were careful about writing “independent” and not “uncorrelated” in the statement above. It is especially this crucial difference we would like to stress in this paper; a difference which lies at the heart of understanding the potential of ART methodology.

The paper is organised as follows. In Section 2 a brief statistical introduction to the notion of copulas is given, especially with applications to ART in mind. Section 3 concentrates on a particular, stylised example from the realm of property insurance. In Section 4 we extrapolate our findings to more general ART products and formulate some conclusions.

2 Measuring dependence through copulas

By now, the history and current development of copula theory is well established. An influential paper importing copulas into mainstream risk management is Embrechts et al. (2002), only published in 2002 though a first version of that paper was available as a preprint from late 1998 onwards. See for instance Embrechts et al. (1999) for a short summary version. The former paper contains numerous references to basic publications in the copula field. For the purpose of the present paper, we highlight some of the basic results, and this mainly through examples. For an in depth discussion, see the above papers and for instance Nelsen (1999) for a textbook treatment. Papers on applications of copula techniques to finance and insurance now appear very regularly; see for instance Schmidt and Ward (2002) for a recent example.

For the purpose of this paper, we concentrate on d 1-period, dependent risks X_1, \dots, X_d , where the dimension $d \geq 2$. In order to price a financial instrument $\Psi(\mathbf{X})$ on $\mathbf{X} = (X_1, \dots, X_d)^T$, ideally one would need full information on the joint probability distribution

$$F_{\mathbf{X}}(\mathbf{x}) = \mathbb{P}(X_1 \leq x_1, \dots, X_d \leq x_d), \quad \mathbf{x} \in \mathbb{R}^d.$$

From this, the distribution function $F_{\Psi(\mathbf{X})}$ of $\Psi(\mathbf{X})$ can be calculated, and hence the position $\Psi(\mathbf{X})$ priced. Since typically we only have partial information on a model for \mathbf{X} , no exact solution can be obtained. Often one knows the distribution functions (or a statistical estimation of) F_{X_1}, \dots, F_{X_d} , the marginal risks, where $F_{X_i}(x) = \mathbb{P}(X_i \leq x)$, $x \in \mathbb{R}$, together with “some information” on the dependence between the risks X_1, \dots, X_d . From now on, we denote $F = F_{\mathbf{X}}$ and $F_i = F_{X_i}$. To avoid technicalities, we assume that the F_i 's are continuous, strictly increasing. From a mathematical point of view, a copula is a distribution function C on the hypercube $[0, 1]^d$ with uniform-(0,1) marginals. A copula provides the natural link between F and (F_1, \dots, F_d) in the following:

$$\begin{aligned} F(\mathbf{x}) &= \mathbb{P}(X_1 \leq x_1, \dots, X_d \leq x_d) \\ &= \mathbb{P}(F_1(X_1) \leq F_1(x_1), \dots, F_d(X_d) \leq F_d(x_d)) \\ &= \mathbb{P}(U_1 \leq F_1(x_1), \dots, U_d \leq F_d(x_d)) \\ &= C(F_1(x_1), \dots, F_d(x_d)). \end{aligned} \tag{1}$$

Here we used that the random variables $F_i(X_i)$ have a uniform-(0,1) distribution, a result used daily for simulation purposes in every bank or insurance

company. The link between X_i and U_i is also referred to as quantile transformation; see Lemma 4.1.9 in Embrechts et al. (1997). One should really write $C = C_F$ in (1), however, we shall drop this F -index. The converse representation to (1) is, for all $u_i \in (0, 1)$, $i = 1, \dots, d$,

$$C(u_1, \dots, u_d) = F(F_1^{-1}(u_1), \dots, F_d^{-1}(u_d)). \quad (2)$$

The formulas (1) and (2) form the basis for all copula modelling. For non-continuous F_i 's one has to be a bit more careful.

Hence C in (1) links the marginals F_i , $i = 1, \dots, d$ to the joint distribution F . Numerous joint models can be considered keeping the F_i 's fixed and varying the copula C . Via (2) we can extract, at least theoretically, the copula function from any joint distribution (i.e. model) F . To start with the latter, suppose for instance that the risk factors X_1, \dots, X_d are jointly normally distributed with covariance matrix Σ ; for convenience we assume that the marginal means are 0 and the variances are 1 so that Σ corresponds to the correlation matrix of \mathbf{X} . Hence for each $i = 1, \dots, d$, $F_i = N$, the standard (one-dimensional) normal distribution function, whereas F corresponds to the above multivariate normal distribution. Hence C in (2) can be calculated numerically leading to the so-called Gaussian copula $C = C_{\Sigma}^{Ga}$. A special choice corresponds to the equicorrelation case where we assume that all correlations ρ_{ij} are equal, to ρ say, in which case we write C_{ρ}^{Ga} . Another important example is the t -copula, extracted via (2) from the multivariate t -distribution on ν degrees of freedom and a positive definite matrix Σ , which for $\mathbf{x} \in \mathbb{R}^d$ has density

$$f_{\mathbf{X}}(\mathbf{x}) = \frac{\Gamma(\frac{\nu+d}{2})}{\Gamma(\frac{\nu}{2}) \sqrt{(\pi\nu)^d |\Sigma|}} \left(1 + \frac{(\mathbf{x} - \boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu})}{\nu} \right)^{-\frac{\nu+d}{2}}. \quad (3)$$

Note that, for $\nu > 2$, the covariance of \mathbf{X} equals $\text{Cov}(\mathbf{X}) = \frac{\nu}{\nu-2} \Sigma$. The normal case above corresponds to the limiting case $\nu = \infty$. Denote the resulting copula by $C_{\Sigma}^{t\nu}$, and $C_{\rho}^{t\nu}$ in the equicorrelation case.

Besides these two natural examples of copulas, one can construct whole families of interesting ones. For instance, for $d = 2$, the following copulas are useful:

- Gumbel copula:

$$C_{\beta}^{Gu}(u, v) = \exp \left\{ - \left[(-\ln u)^{\beta} + (-\ln v)^{\beta} \right]^{1/\beta} \right\}, \quad \beta \in [1, \infty),$$

- Frank copula:

$$C_{\beta}^{Fr}(u, v) = -\frac{1}{\beta} \ln \left(1 + \frac{(e^{-\beta u} - 1)(e^{-\beta v} - 1)}{e^{-\beta} - 1} \right), \quad \beta \in \mathbb{R} \setminus \{0\},$$

- Clayton copula:

$$C_{\beta}^{Cl}(u, v) = \max\left(\left(u^{-\beta} + v^{-\beta} - 1\right)^{-1/\beta}, 0\right), \quad \beta \in [-1, \infty) \setminus \{0\},$$

- Farlie–Gumbel–Morgenstern copula:

$$C_{\beta}^{FGM}(u, v) = uv + \beta uv(1 - u)(1 - v), \quad \beta \in [-1, 1].$$

For generalisations to $d > 2$, see for instance Nelsen (1999) or Embrechts et al. (2000). These publications also explain how to prove that the above functions are indeed copulas.

Below we illustrate for the more applied reader some of the key facts of copula modelling. First of all, we would like to stress that the key question remains: “How to find the copula that best models the dependence structure of real multivariate data.” Though we will partly address this question, it is important to state already at the beginning that giving an answer is essentially as difficult as estimating the joint distribution F in the first place; hence for d large difficult, if not impossible. At the moment, the main virtues of copula modelling are i) pedagogic: opening up the mind of the user to think about dependence beyond linear correlation (see Embrechts et al. (2002) for several examples on this), and ii) stress testing (typically using simulation) the pricing of financial (one–period) products on multivariate financial or insurance underlyings (typical examples in finance include basket options and credit derivatives; in insurance, multi–line products).

We have decided to illustrate below some of the key features of copulas using simulation examples. Figure 1 contains three models each having standard normal marginals ($F_1 = F_2 = N$), illustrated by the marginal histograms. Moreover all three models have the same linear correlation of 50%, but the copula changes. In Figure 1.a $C = C_{0.5}^{Ga}$, 1.b $C = C_{0.5}^{t_4}$, 1.c $C = C_{\beta}^{Gu}$ where β is chosen in such a way that the resulting bivariate distribution has 50% correlation. In each plot we have simulated 1000 points and empirically estimated the linear and Kendall τ rank correlation; for a definition of the latter, see Embrechts et al. (2002). A key feature to be observed is the change from a perfectly elliptical cloud with few joint extremes (beyond 2, say) in Figure 1.a, over a two–sided elongated cloud with more extremes both NE and SW in Figure 1.b, to a teardrop shape in Figure 1.c with peak (joint extremes) mainly in the NE corner. It is to be stressed once more that these important model differences cannot be explained on the basis of marginals and linear correlations alone, indeed these are identical across the figures! To further stress this point, look at

Figure 2. The lines drawn in each plot correspond to univariate quantiles q of the marginal distribution so that $\mathbb{P}(X_1 > q, X_2 > q) = 0.99$, added only for illustrative purposes in order to indicate a “danger zone” in the joint distribution. Here we again have $F_1 = F_2 = N$, $\rho = 50\%$ in Figure 2.a (the bivariate normal case). Note the less pronounced extreme joint moves (NE, SW). Joint extreme moves appear much more pronounced in Figure 2.d,

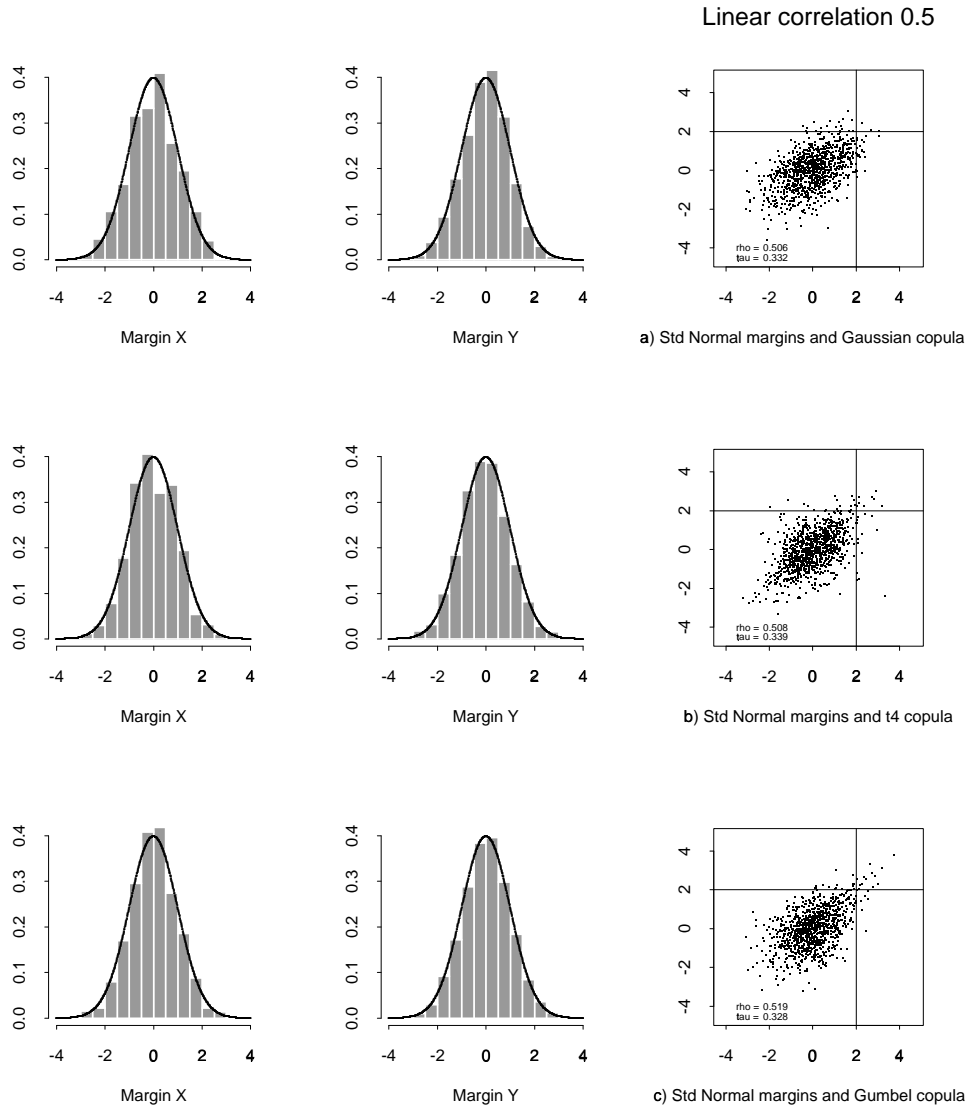


Figure 1: *Bivariate distributions with fixed correlation and different copulas.*

the case of a bivariate t (see (3)) with $\nu = 4$, $\rho = 50\%$. In the latter case, we have that $F_1 = F_2 = t_4$. Moving from Figure 2.a to 2.c we keep $F_1 = F_2 = N$ but (using (1)) superimpose the copula $C_{0.5}^{t_4}$ on F_1, F_2 . We now see points clustering more into the NE and SW corners. Between the Figures 2.a and 2.b we change the marginals (from N to t_4) but keep the copula fixed

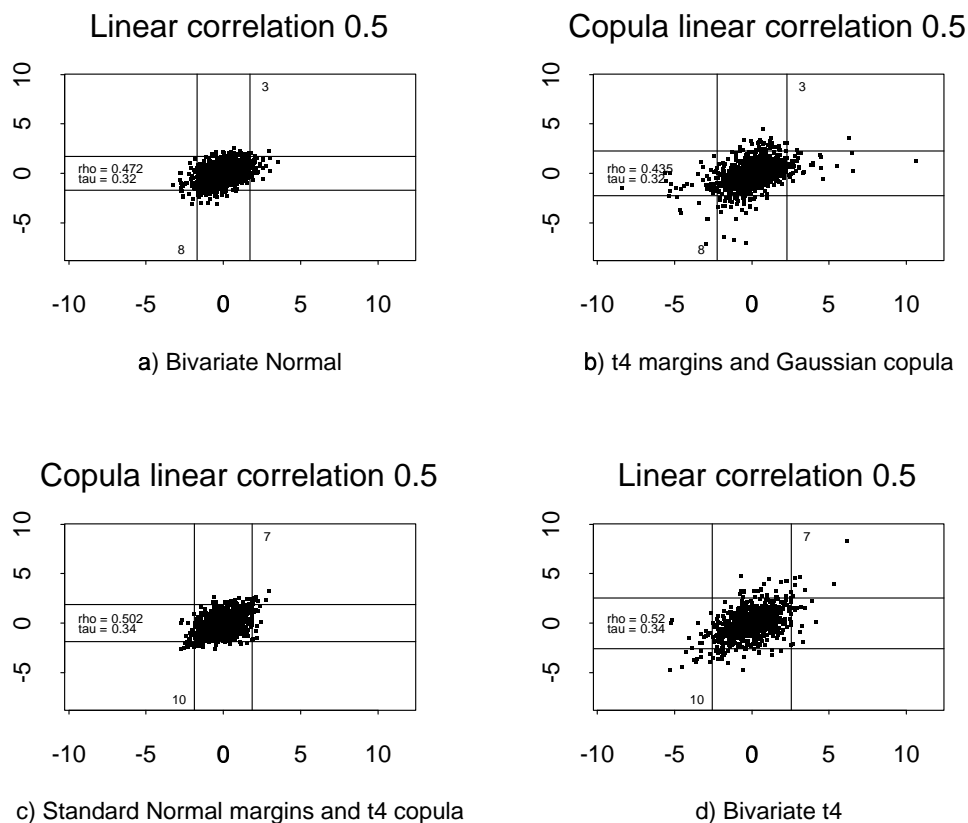


Figure 2: *Four basic models around the bivariate normal and the bivariate t_4 distribution.*

($C_{0.5}^{Gg}$). Making the marginal tails much longer (from N to t_4) but keeping the copula Gaussian does not yield pronounced joint extremes. From these figures we get the feeling that joint extremes are more a copula property than a marginal distribution one. This crucial fact is indeed true and can be shown mathematically; see for instance Embrechts et al. (2002). Figures 2.a and 2.d correspond to the class of so-called elliptical distributions. In the Figures 2.b and 2.c, the resulting models fail to have correlation 50%, though the estimated values show that the deviation is fairly small. We could have calibrated these models differently by changing $\nu = 4$ to achieve a theoretical value of 50% for ρ . It is especially this flexibility of modelling separately the marginals (F_1, F_2) and the copula (C) that will allow us to come up with models which reflect more closely the qualitative behaviour of multivariate loss data. One could call this approach “copula engineering”. In the next section we give an illustrative example of this technique based on some real insurance loss data.

3 An example: the Danish fire data

For a one-dimensional discussion of these data, see Embrechts et al. (1997), Example 6.2.9 and further. See also Rytgaard (1996), from whom we obtained the data, and McNeil (1997) for an early extreme value analysis. For our purposes important to realise is that the data ($n = 2493$) is three-dimensional ($d = 3$). Indeed, the losses to industrial dwellings consist of $X_1 =$ loss to buildings, $X_2 =$ loss to contents, $X_3 =$ loss to profits. For the data which have for each component a strictly positive loss amount (resulting in a reduced sample size of $n = 517$), we have plotted the histograms for the log-transformed variables; see Figure 3. We denote $Y_i = \log(X_i)$, $i = 1, 2, 3$.




Figure 3: *Marginal histograms of the log-transformed Danish fire data.*

From Figure 3.a we clearly see heavy-tailedness (beyond lognormality) in the marginal data, further exemplified in Figure 4.




Figure 4: *QQ-plots of log-transformed data against normality.*

The data are clearly dependent, as can be seen from the three-dimensional (Figure 5) and two-dimensional (Figure 6) scatterplots.

Suppose, as an example, that one would be interested in a bivariate model for log-contents (Y_2) versus log-profits (Y_3) (see Figure 6.c). A very important property of a copula is that it is invariant under continuous and strictly increasing transformations (like the logarithm) of the underlying

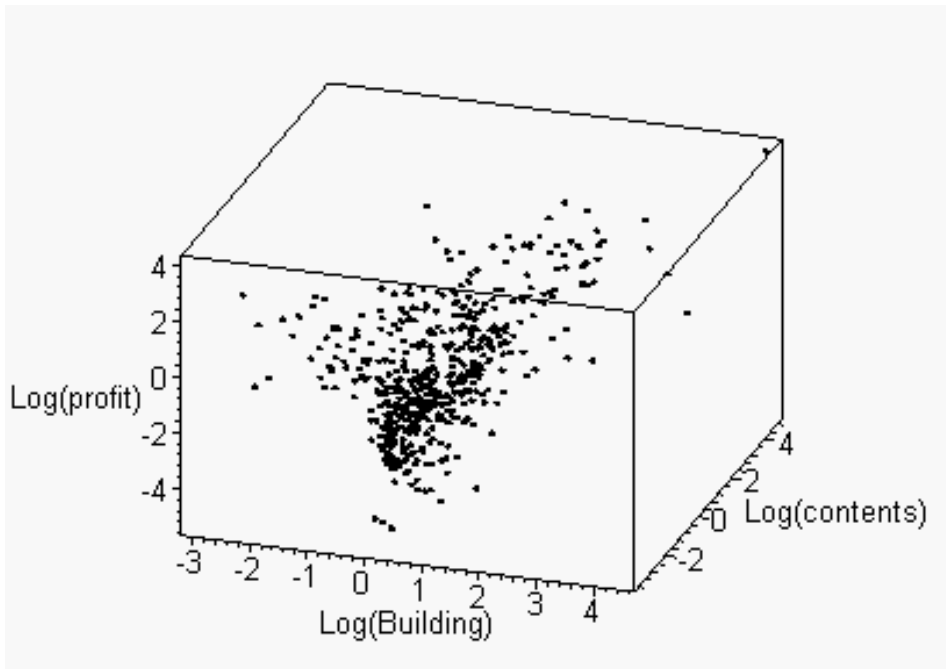


Figure 5: *Three-dimensional scatterplot of the log-transformed Danish fire data.*

variables; see for instance Proposition 2 in Embrechts et al. (2002). Linear correlation does not have this invariance property; for this, and further shortcomings of linear correlation, see Embrechts et al. (2002), Section 3.2. Hence it makes no difference to fit the copula to the data (X_i) or to the logarithms of the data (Y_i) . One possibility for obtaining a statistical model would be by using a so-called pseudo-likelihood approach as discussed in Genest et al. (1995). In this method one replaces the unknown copula C by a parametric copula family C_θ (the Gumbel family, say) and estimates the marginal distributions (here F_2, F_3) empirically. An optimal choice among

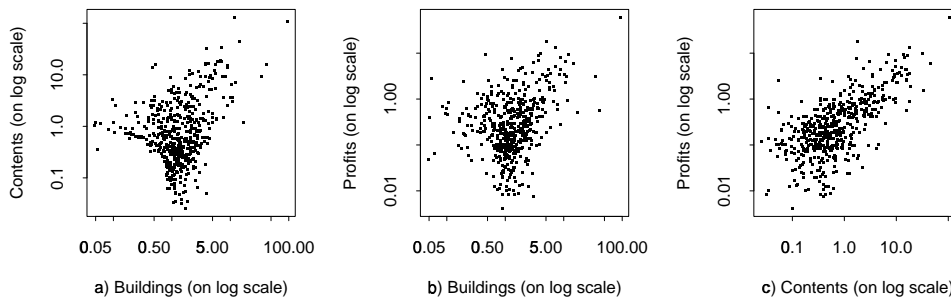


Figure 6: *Bivariate scatterplots of the log-transformed Danish fire data.*

several competing parametric models for C can for instance be obtained

through the minimisation of some information criterion, like Akaike’s AIC. Of course, several parametric or semi-parametric models for the marginal fitting may also be used, possibly based on extreme value theory; see for instance McNeil and Frey (2000).

By way of example, below we have used the pseudo-likelihood approach to fit some copulas to the (X_2, X_3) data. The results are summarised in

Copula family	Parameter space	$\hat{\beta}$		AIC
Gumbel	$\beta \geq 1$	1.8569		-319.51
Clayton	$\beta \in [-1, \infty) \setminus \{0\}$	1.0001		-107.73
Frank	$\beta \in \mathbb{R} \setminus \{0\}$	5.0008		-262.75
F-G-M	$\beta \in [-1, 1]$	0.9996		-165.84
		$\hat{\rho}$	$\hat{\nu}$	
t	$\rho \in [-1, 1], \nu > 0$	0.6456	9.1348	-268.55
Gaussian	$\rho \in [-1, 1]$	0.6323	-	-262.06

Table 1: Parameter estimates for various copula models for (X_2, X_3) .

Table 1. The fit based on the Gumbel copula has the lowest AIC value, so according to this criterion should be chosen. The Clayton family performs worst. The Gaussian and the Frank copulas perform similarly according to the AIC. The Farlie-Gumbel-Morgenstern family, one knows, is not appropriate to fit data exhibiting strong dependence as is this case. This is confirmed by the estimate close to the β -parameter boundary $\beta = 1$. Further diagnostic checks could be added.

We can use these estimates to price positions $\Psi(X_2, X_3)$ based on the variables (X_2, X_3) . Four payoff functions are considered in a simulation study using the six fitted models together with the model assuming independence. The payoff functions chosen are the following:

- $\Psi_1(X_2, X_3) = X_2 + X_3$,
- $\Psi_2(X_2, X_3) = (X_2 + X_3 - 10)_+$,
- $\Psi_3(X_2, X_3) = (X_2 + X_3) \cdot \mathbf{1}_{\{X_2 > 5, X_3 > 5\}}$,
- $\Psi_4(X_2, X_3) = \mathbf{1}_{\{X_2 > 5, X_3 > 5\}}$.

Ψ_1 corresponds to the combined loss, whereas Ψ_2 is like a stop-loss treaty on Ψ_1 . The position Ψ_4 can be viewed as a digital, i.e. taking the value 1 whenever both X_2 and X_3 are larger than 5, otherwise the value 0 is obtained. Position Ψ_3 pays out the combined loss Ψ_1 only if triggered by Ψ_4 .

Monte-Carlo estimates of the fair premium $\mathbb{E}(\Psi_i(X_2, X_3))$ under each of the seven fitted models are obtained through 500’000 simulations and are

given in Table 2; see the $\hat{\mathbb{E}}_{\hat{F}}(\Psi)$ column. A combination of the estimates for the mean ($\hat{\mathbb{E}}_{\hat{F}}(\Psi)$) and the standard deviation ($\hat{\sigma}_{\hat{F}}(\Psi)$) for each position can be used to construct an actuarial (standard deviation) premium principle.

$$\Psi_1(X_2, X_3) = X_2 + X_3$$

Copula model	$\hat{\mathbb{E}}_{\hat{F}}(\Psi)$	$\widehat{s.e.}(\hat{\mathbb{E}}_{\hat{F}}(\Psi))$	$\hat{\sigma}_{\hat{F}}(\Psi)$	$\widehat{s.e.}(\hat{\mathbb{E}}_n(\Psi))$	$\widehat{s.e.}(\hat{\sigma}_n(\Psi))$
Gumbel	3.231	0.015	10.772	0.458	3.060
<i>t</i>	3.224	0.013	10.009	0.445	2.768
Frank	3.205	0.013	9.188	0.409	2.235
Gaussian	3.232	0.014	9.853	0.424	2.516
FGM	3.245	0.013	9.120	0.398	2.254
Clayton	3.243	0.013	9.137	0.409	2.262
Independence	3.234	0.013	8.985	0.397	2.234

$$\Psi_2(X_2, X_3) = (X_2 + X_3 - 10)_+$$

Copula model	$\hat{\mathbb{E}}_{\hat{F}}(\Psi)$	$\widehat{s.e.}(\hat{\mathbb{E}}_{\hat{F}}(\Psi))$	$\hat{\sigma}_{\hat{F}}(\Psi)$	$\widehat{s.e.}(\hat{\mathbb{E}}_n(\Psi))$	$\widehat{s.e.}(\hat{\sigma}_n(\Psi))$
Gumbel	1.134	0.013	9.525	0.406	3.240
<i>t</i>	1.066	0.012	8.711	0.389	2.927
Frank	1.001	0.011	7.801	0.346	2.368
Gaussian	1.072	0.012	8.522	0.367	2.655
FGM	0.970	0.011	7.773	0.346	2.382
Clayton	0.970	0.011	7.789	0.349	2.388
Independence	0.927	0.011	7.686	0.340	2.350

$$\Psi_3(X_2, X_3) = (X_2 + X_3) \cdot \mathbf{1}_{\{X_2 > 5, X_3 > 5\}}$$

Copula model	$\hat{\mathbb{E}}_{\hat{F}}(\Psi)$	$\widehat{s.e.}(\hat{\mathbb{E}}_{\hat{F}}(\Psi))$	$\hat{\sigma}_{\hat{F}}(\Psi)$	$\widehat{s.e.}(\hat{\mathbb{E}}_n(\Psi))$	$\widehat{s.e.}(\hat{\sigma}_n(\Psi))$
Gumbel	1.118	0.014	10.382	0.456	3.518
<i>t</i>	0.727	0.011	8.067	0.348	3.187
Frank	0.303	0.006	4.201	0.196	2.183
Gaussian	0.657	0.010	7.336	0.322	2.943
FGM	0.142	0.004	2.793	0.123	1.791
Clayton	0.146	0.004	2.805	0.128	1.794
Independence	0.077	0.003	2.128	0.097	1.643

$$\Psi_4(X_2, X_3) = \mathbf{1}_{\{X_2 > 5, X_3 > 5\}}$$

Copula model	$\hat{\mathbb{E}}_{\hat{F}}(\Psi)$	$\widehat{s.e.}(\hat{\mathbb{E}}_{\hat{F}}(\Psi))$	$\hat{\sigma}_{\hat{F}}(\Psi)$	$\widehat{s.e.}(\hat{\mathbb{E}}_n(\Psi))$	$\widehat{s.e.}(\hat{\sigma}_n(\Psi))$
Gumbel	0.0227	0.0133	0.1489	0.0064	0.0209
<i>t</i>	0.0160	0.0121	0.1255	0.0053	0.0209
Frank	0.0095	0.0001	0.0968	0.0043	0.0229
Gaussian	0.0015	0.0119	0.1229	0.0054	0.0218
FGM	0.0048	0.0108	0.0689	0.0031	0.0266
Clayton	0.0049	0.0108	0.0700	0.0031	0.0268
Independence	0.0025	0.0107	0.0023	0.0502	0.0289

Table 2: Monte–Carlo analysis for four positions on the (X_2, X_3) data from seven different models.

Note that the values in Table 2 are derived from specific models fitted to the data. In Table 3 we have summarised some empirical estimates of the most important quantities, hence no model fitting is taking place here. The various standard error (*s.e.*) estimates listed in Table 2 can be used to further discuss model risk and fitting adequacy (goodness–of–fit); we will

Payoff	Sample statistics		
	$\hat{\mathbb{E}}_n(\Psi)$	$\widehat{s.e.}(\hat{\mathbb{E}}_n(\Psi))$	$\hat{\sigma}_n(\Psi)$
$\Psi_1(X_2, X_3)$	3.233	0.467	10.624
$\Psi_2(X_2, X_3)$	1.149	0.411	9.335
$\Psi_3(X_2, X_3)$	0.819	0.367	8.338
$\Psi_4(X_2, X_3)$	0.0213	0.0064	0.1444

Table 3: Empirical estimates for four positions on the (X_2, X_3) data.

refrain from commenting further on these values but have included them for completeness. Recall that the Gumbel model was chosen as the best fitting model (accounting to AIC); this choice is also largely confirmed when we compare the estimated mean values and standard deviations in Tables 2 and 3. Also note the, in some cases large, differences between the estimated fair premiums. A careful modelling of the dependence can make a huge difference in the quoted prices. On the basis of this fairly preliminary analysis, we would choose the Gumbel model for the copula superimposed via (1) on the marginal empirical distributions, i.e.

$$\hat{F}(x_2, x_3) = C_{1.8569}^{Gu}(\hat{F}_{2,n}(x_1), \hat{F}_{3,n}(x_2)).$$

Similar analysis can be performed for the other bivariate claims, as well as for the trivariate data. We only make some comments. For the pair (X_1, X_2) we notice the special wedge like behaviour near zero in Figure 6.a. In order to come up with a good statistical model in this case, we would have to consider so-called censoring techniques as discussed in Shih and Louis (1995). In the case of (X_1, X_3) we notice a clear diagonal asymmetry in the scatterplot of Figure 6.b implying that we have to look for non-exchangeable copulas; see Genest et al. (1998) and Joe (1997). For a good trivariate model one would have to combine the above approaches. We refrain from entering into more details here.

4 The importance of copula techniques for ART

How do the above techniques reflect on the modelling of ART products? Traditionally, insurance companies have managed each risk separately, e.g. by buying reinsurance for a certain line of business (insurance risk factor), or by setting up hedges for some foreign currency (financial risk factor). The quantitative analysis of such isolated risk management tools did not (or only rarely) require a deep understanding of dependence between the various risk factors.

While these isolated risk management tools are relatively easy to handle and easy to analyse, they do not necessarily offer the most cost-effective

protection if one looks at the company as a whole. The fact that a company depends on various stochastic risk factors creates the possibility for diversification, i.e. the use of offsetting effects between the risk factors. This fact has been well-known for a long time, e.g. in insurance through the work of de Finetti in 1942 or Borch's Theorem in the late sixties (see Bühlmann (1970) for both) and in finance through Markowitz's portfolio theory from 1952 (see e.g. Brealy and Myers (2000)). But only recently did insurance companies start to capitalise on this diversification potential, driven by factors such as

- increased pressure from the investors for better overall results and, thus, better returns on the risk capital that they provide to the insurers,
- lack of capacity in some reinsurance markets, and hence the need to use available capacity more efficiently or to open up new capacity.

As a matter of fact, an insurance company needs not to protect itself against adverse outcomes in single risk factors; it is basically sufficient to protect the company's bottom line while taking into account the diversification potential emanating from the variety of the risk factors¹. The new types of risk transfer solutions that were created in response to these new approaches to risk management are now known under the name Alternative Risk Transfer (ART). An easily accessible account on the trend towards the integrated management and transfer of risks as well as descriptions of the most important types of ART products is given in Shimpi (1999).

Practical ART solutions tend to be highly customised and rather complicated, partly due to regulatory, accounting and tax implications, and it is well beyond the scope of our paper to treat detailed examples. Instead, we investigate some simple, stylised classes of ART solutions in order to make our point that proper understanding and modelling of dependence structures among the risk factors is essential for any related quantitative analysis.

- *Finite Risk Covers* are normally aimed at spreading the risk of one (insurance) risk factor over time, which is outside the scope of this paper. However, given the long time horizons of such deals, it is fairly commonplace to incorporate investment income into the pricing, which - in turn - requires the understanding of possible dependence between investment return and insured risk.

¹We tacitly assume here that corporate governance structures are set up in such a way that they do not penalise adverse results of a single line of business if they are due to an integrated risk management structure that focuses on protection of the overall bottom line only.

- *Multi-Line Products* cover the total risk of several risk factors, with whole account stop-loss treaties being a simple and popular example. Understanding the dependence structure of the included risk factors is therefore essential for understanding the behaviour of the total risk. Multi-line deals are often underwritten for several years, which makes it sensible to incorporate investment income, including - of course - possible dependences with the other risk factors.
- *Multi-Trigger Products* provide payments only if several conditions are fulfilled simultaneously; with knock-in options being a popular product from finance. Although pure multi-trigger products are only rarely underwritten in insurance, multi-trigger structures are regularly part of other products. Multi-trigger structures allow to trigger payoffs only if they are actually necessary, e.g. some large insurance loss is only paid if the company has simultaneously suffered losses on its asset side, indicated by stock and bond indices going below some threshold. This reduces the probability of payoffs and thus the premium for the cover. Almost needless to say that sensible design and pricing of such products is only possible if the dependence between the triggering risk factors is well understood.
- *Securitization Products* (e.g. CAT bonds) can be based on multiple lines, or they can have multiple triggers, e.g. in order to avoid moral hazard. In any case, dependence modelling is important for yet another reason: as already stated in Section 1, these products are sold to investors in the capital markets with the marketing argument that they have low, or even negative, correlation with the usual investment assets and thus provide good diversification of investment portfolios.

Recall that $\mathbf{X} = (X_1, \dots, X_d)^T$ denotes the vector of random risk factors and $\mathbf{a} = (a_1, \dots, a_d)^T \in (\mathbb{R} \cup \{\pm\infty\})^d$ some constant vector. We can easily formalise the payoffs of the above-mentioned types of stylised ART products as either being a function of the aggregate sum, i.e.

$$\Psi(\mathbf{X}) = f\left(\sum_{i=1}^d g(X_i)\right)$$

or as being contingent on the risk factors assuming certain values, i.e.

$$\Psi(\mathbf{X}) = c \cdot \mathbf{1}_{\{X_1 > a_1, \dots, X_d > a_d\}}$$

or as being combinations thereof, e.g.

$$\Psi(\mathbf{X}) = f\left(\sum_{i=1}^d g(X_i)\right) \cdot \mathbf{1}_{\{X_1 > a_1, \dots, X_d > a_d\}}$$

where, sometimes, we are only interested in conditional distributions, e.g. the distribution of

$$f(X_1, \dots, X_d) \text{ given that } X_1 > a_1, \dots, X_d > a_d.$$

Note that the payoff functions Ψ_1, \dots, Ψ_d from Section 3 all belong to one of the above classes of stylised payoff functions. It (almost) goes without saying that the distribution of the payoff $\Psi(\mathbf{X})$ depends in all cases on the joint distribution F of the risk factors, i.e. on the marginal distributions of the single risk factors X_1, \dots, X_d as well as on their dependence structure (copula).

Remember moreover that the risk factors X_1, \dots, X_d typically include pure insurance risks (downside only) as well as financial risks (double-sided). In ART set-ups, we are therefore not likely to face risk vectors \mathbf{X} for which an element of the class of elliptical distributions is a sensible model. Elliptical distributions can be thought of roughly as a generalisation of the multivariate Gaussian distribution. However, in these cases where the joint distribution of the risk factors \mathbf{X} is non-elliptical, the usual linear correlation is not a viable means for describing the dependence structure between the risk factors. See Embrechts et al. (2002) for more details on these issues. Therefore, we have to revert to more sophisticated concepts for dependence modelling, as for instance the copula approach described in Sections 2 and 3.

Another aspect is that, in many cases, ART products are designed in such a way that they are hit only with relatively low probability (e.g. CAT bonds). Hence, whatever the exact structure of such products is, the interest often lies on “rare” or “extreme” events of the joint distribution of the risk factors. Therefore, dependence modelling must be particularly accurate in the related regions of the joint distribution function F . This is the more so in view of the observation that dependence structures that prevail for the “usual” course of events of the risk factors tend to change under more extreme events. Think of two risk factors X_1 and X_2 that are almost uncorrelated for high-probability events, but that tend to assume particularly high values jointly, as e.g. in Figure 1.c. For instance a small fire destroys some furniture in a commercial building, but does not cause a business interruption. A large fire, however, generates a high claim from destroyed

furniture and, moreover, is likely to cause also a major business interruption with an associated large claim. A worked-out example with data from finance as well as a theoretical result on copula convergence for tail events given in Juri and Wüthrich (2001) provides further evidence for this.

Classical portfolio theory along the lines of Markowitz and CAPM, as described for instance in Brealy and Myers (2000) is closely related to the use of linear correlation for the measurement and modelling of dependence between the considered classes of assets. Therefore, the above-stated arguments suggest that it may not be optimal to use classical portfolio theory to justify the usefulness of insurance-linked securities (like CAT bonds) to the investor, since the joint distribution of insurance risks and financial risks is not likely to be elliptical. A possible alternative approach based on VaR and copulas for the dependence modelling is given in Embrechts et al. (2001).

We now turn back to our data example in order to get an idea of the impact of dependence modelling on the distributions of some stylised ART payoffs and hence on the pricing of such products. Putting the results of the computations for the different models and payoff function examples from Section 3 together, we obtain the following results for the expected values of the payoffs:

Model	AIC	$\hat{\mathbb{E}}(\Psi_1)$	$\hat{\mathbb{E}}(\Psi_2)$	$\hat{\mathbb{E}}(\Psi_3)$	$\hat{\mathbb{E}}(\Psi_4)$
Gumbel	-319.51	3.231	1.134	1.118	0.0227
t	-268.55	3.224	1.066	0.727	0.0160
Frank	-262.75	3.205	1.001	0.303	0.0095
Gaussian	-262.06	3.232	1.072	0.657	0.0015
FGM	-165.84	3.245	0.970	0.142	0.0048
Clayton	-107.73	3.243	0.970	0.146	0.0049
Independence	-	3.234	0.927	0.077	0.0025

and for the standard deviations of the payoffs:

Model	AIC	$\hat{\sigma}(\Psi_1)$	$\hat{\sigma}(\Psi_2)$	$\hat{\sigma}(\Psi_3)$	$\hat{\sigma}(\Psi_4)$
Gumbel	-319.51	10.772	9.525	10.382	0.1489
t	-268.55	10.009	8.711	8.067	0.1255
Frank	-262.75	9.188	7.801	4.201	0.0968
Gaussian	-262.06	9.853	8.522	7.336	0.1229
FGM	-165.84	9.120	7.773	2.793	0.0689
Clayton	-107.73	9.137	7.879	2.805	0.0700
Independence	-	8.985	7.686	2.128	0.0023

Recall that all models use the same marginal distributions (i.e. any difference in the result is due to the choice of the different dependence structures). The AIC is a measure for the goodness-of-fit, where lower values indicate a better fit, i.e. the results are ordered in descending order of

goodness-of-fit (except for the independence assumption). We can observe the following:

- There are indeed significant differences between the estimates generated under the various dependence models. In particular, there are also significant differences between the best-fitting models and the independence assumption. In particular, the estimates under the independence assumption usually correspond well with estimates from models with poor goodness-of-fit values.
- The differences between the estimates usually correspond well with the differences in the goodness-of-fit. In particular, models with similar goodness-of-fit value usually produce similar estimates.

Figure 7 shows the estimated cumulated distribution functions of the payoff $\Psi_1(\mathbf{X}) = X_2 + X_3$ for two different dependence models and under the assumption of independence. From this figure it can also be seen that whether or not to account for dependence, and if so, how, can make a significant difference in the resulting profit-and-loss distribution.




Figure 7: *Estimated distribution functions of $\Psi_1(\mathbf{X})$ resulting from different model assumptions.*

Given that the pricing of insurance and ART products relies heavily on the expectation of the payoffs and their dispersion, we conclude that the

accurate measurement and modelling of dependence between risk factors is a crucial ingredient for being able to determine effective product structures and viable prices. In this paper we have shown, through some examples, that one has to go beyond the notion of linear correlation in order to capture the true dependence underlying typical insurance (and indeed finance) data. The notion of copula can help to achieve a better understanding of the role that dependence plays in the realm of ART. We would like to stress that our contribution just focused on some of the underlying issues, no doubt considerable additional work in this area is needed.

References

- Brealy, R. A. and Myers, S. C. (2000). *Principles of Corporate Finance (6th Edition)*. McGraw-Hill, New York.
- Bühlmann, H. (1970). *Mathematical Methods in Risk Theory*. Springer, Berlin.
- Embrechts, P., Hoeing, A., and Juri, A. (2001). Using Copulae to bound the Value-at-Risk for functions of dependent risks. Working paper, ETH Zürich.
- Embrechts, P., Klüppelberg, C., and Mikosch, T. (1997). *Modelling Extremal Events for Insurance and Finance*. Springer, Berlin.
- Embrechts, P., Lindskog, F., and McNeil, A. J. (2000). Modelling dependence with copulas and applications to risk management. To appear in *Heavy-Tailed Distributions in Finance (2002)*(S.T. Rachev, Ed.) North Holland, Amsterdam.
- Embrechts, P., McNeil, A. J., and Straumann, D. (1999). Correlation: pitfalls and alternatives. *Risk Magazine*, 5:69–71.
- Embrechts, P., McNeil, A. J., and Straumann, D. (2002). Correlation and dependence in risk management: Properties and pitfalls. In: *Risk Management: Value at Risk and Beyond* (M. Dempster, Ed.) Cambridge University Press, Cambridge, 176–223.
- Genest, C., Ghoudi, C. K., and Rivest, L.-P. (1998). Commentaire sur l'article de E. W. Frees et E. A. Valdez intitulé "Understanding relationships using copulas". *North American Actuarial Journal*, 2:143–149.

- Genest, C., Ghoudi, K., and Rivest, L.-P. (1995). A semiparametric estimation procedure of dependence parameters in multivariate families of distributions. *Biometrika*, 82(3):543–552.
- Gibson, R., Ed. (2000). *Model Risk, Concepts, Calibration and Pricing*. Risk Waters Group, London.
- Gisler, A. and Frost, P. (1999). An addendum and a short comment on the paper from U. Schmock, “Estimating the value of the WinCAT coupons of the Winterthur Insurance convertible bond: a study of the model risk”. *Astin Bulletin*, 29(1):165–171.
- Joe, H. (1997). *Multivariate Models and Dependence Concepts*. Chapman & Hall, London.
- Juri, A. and Wüthrich, M. (2001). Copula convergence theorems for tail events. Working paper, ETH Zürich.
- McNeil, A. J. (1997). Estimating the tails of loss severity distributions using extreme value theory. *Astin Bulletin*, 27(1):117–137.
- McNeil, A. J. and Frey, R. (2000). Estimation of tail-related risk measures for heterocedastic financial time series: an extreme value approach. *Journal of Empirical Finance*, 7:271–300.
- Nelsen, R. B. (1999). *An Introduction to Copulas*. Springer, New York.
- Rytgaard, M. (1996). Simulation experiments on the mean residual life function $m(x)$. In: *Proceedings of the XXVII ASTIN Colloquium, Copenhagen, Denmark*, 1:59–81.
- Schmidt, W. and Ward, I. (2002). Pricing default baskets. *Risk*, 1:111–114.
- Schmock, U. (1999). Estimating the value of the WinCAT coupons of the Winterthur Insurance convertible bond: a study of the model risk. *Astin Bulletin*, 29(1):101–163.
- Shih, J. H. and Louis, T. A. (1995). Inferences on the association parameter in copula models for bivariate survival data. *Biometrics*, 51:1384–1399.
- Shimpi, P., Ed. (1999). *Integrating Corporate Risk Management*. Swiss Re New Markets, New York.