

Total loss estimation using copula-based regression models[☆]

Nicole Krämer^a, Eike C. Brechmann^{a,1}, Daniel Silvestrini^a, Claudia Czado^a

^a*TU München, Department of Mathematical Statistics, Parkring 13, 85748 Garching, Germany*

Abstract

We present a joint copula-based model for insurance claims and sizes. It uses bivariate copulae to accommodate for the dependence between these quantities. We derive the general distribution of the policy loss without the restrictive assumption of independence. We illustrate that this distribution tends to be skewed and multi-modal, and that an independence assumption can lead to substantial bias in the estimation of the policy loss. Further, we extend our framework to regression models by combining marginal generalized linear models with a copula. We show that this approach leads to a flexible class of models, and that the parameters can be estimated efficiently using maximum-likelihood. We propose a test procedure for the selection of the optimal copula family. The usefulness of our approach is illustrated in a simulation study and in an analysis of car insurance policies.

Keywords: dependence modeling, generalized linear model, number of claims, claim size, policy loss

[☆]All numerical computations for this paper were performed on a Linux cluster supported by DFG grant INST 95/919-1 FUGG.

¹The second author gratefully acknowledges the support of the TUM Graduate School's International School of Applied Mathematics as well as of Allianz Deutschland AG.

1. Introduction

Estimating the total loss of an insurance portfolio is crucial for many actuarial decisions, e.g. for pricing of insurance contracts and for the calculation of premiums. A very common approach, based on the compound model by Lundberg (1903), models the average claim size and the number of claims independently, and then defines the loss as the product of these two quantities. However, the assumption of independence can be too restrictive and lead to a systematic over- or under-estimation of the policy loss. Evidently, this effects the accuracy of the estimation of the portfolio loss.

We therefore propose a joint model that explicitly allows a dependency between average claim sizes and number of claims. This is achieved by combining marginal distributions for claim frequency and severity with families of bivariate copulae. A main contribution of this paper is the derivation of the distribution of the loss of an insurance policy. We illustrate that the distribution is often very skewed, and that – depending on the model parameters – the distribution is multi-modal. Based on this distribution, we can estimate the expected policy loss and its quantiles. We show that the distribution, and in particular its mean, depends strongly on the degree of dependence. This underpins the usefulness of our copula-based model.

Dependence modeling using copulae has become very popular the last years (see the standard reference books by Joe (1997) and Nelsen (2006)) and was introduced to actuarial mathematics by Frees and Valdez (1998). Since then, copulae have been used, e.g., for the modeling of bivariate loss distributions by Klugman and Parsa (1999) and of dependencies between

loss triangles by De Jong (2009), as well as for risk management (see McNeil et al. (2005)).

It is common practice to model average claim sizes and number of claims in terms of a set of covariates as e.g. gender or age (see, e.g., Haberman and Renshaw (1996) for an overview). Typically, claim frequency and severity are however modeled separately under the independence assumption of Lundberg (1903). Gschlöbl and Czado (2007) therefore included the number of claims as a covariate into the model for average claim size. To allow for more flexibility and generality in the type of dependence, we extend our copula-based model to regression models by combining generalized linear models for the two marginal regression models with copula families. This is an extension of a recent approach by Czado et al. (2012) and De Leon and Wu (2011) who only consider a Gauss copula based on work by Song (2000, 2007) and Song et al. (2009).

In our general copula-based regression approach, the model parameters can be estimated efficiently using maximum-likelihood techniques. Further, we provide asymptotic confidence intervals that allow us to quantify the uncertainty of our estimates. For the selection of the copula family, we propose the likelihood-ratio test by Vuong (1989).

In an extensive simulation study, we show that the incorporation of a copula allows a more precise estimation of the individual policy losses, which in turn leads to a more reliable estimation of the total loss. These results are confirmed in a case study on car insurance policies. All statistical methods presented here are implemented in the R package 'CopulaRegression' (Krämer and Silvestrini, 2012).

2. Bivariate copulae for continuous-discrete data

2.1. Background: bivariate copulae

A bivariate copula $C : [0, 1] \times [0, 1] \rightarrow [0, 1]$ is a bivariate cumulative distribution function on $[0, 1] \times [0, 1]$ with uniformly distributed margins. The importance of copulae is underpinned by Sklar's Theorem (1959). In the bivariate case, it states that for every joint distribution function $F_{X,Y}$ of a bivariate random variable (X, Y) with univariate marginal distribution functions F_X and F_Y , there exists a bivariate copula C such that

$$F_{X,Y}(x, y) = C(F_X(x), F_Y(y)). \quad (1)$$

If X and Y are continuous random variables, the copula C is unique. Conversely, if C is a copula, Equation (1) defines a bivariate distribution with marginal distribution functions F_X and F_Y . This allows us to model the marginal distributions and the joint dependence separately, as we can define the copula C independently of the marginal distributions.

Copulae are invariant under monotone transformations of the marginal distributions. Therefore, instead of the correlation coefficient – which measures linear associations – monotone association measures are used. A very common choice is Kendall's τ ,

$$\tau := 4 \int_{[0,1]^2} C(u, v) dC(u, v) - 1 \in [-1, 1].$$

In this paper, we study copula-based models for a pair of continuous-discrete random variables. We denote the continuous random variable by X , and the discrete random variable by Y . We assume that Y takes values in $1, 2, \dots$

Their joint distribution is defined by a parametric copula $C(\cdot, \cdot | \theta)$ that depends on a parameter θ , i.e. the joint distribution is given by

$$F_{X,Y|\theta}(x, y) = C(F_X(x), F_Y(y) | \theta).$$

We focus on four families of parametric bivariate copulae, namely the Clayton, Gumbel, Frank and Gauss copulae. Each family depends on a copula parameter θ . These parameters can be expressed in terms of Kendall's τ . The definitions of the copula families and their relationship to Kendall's τ are provided in Appendix A. We note that the Clayton copula is only defined for positive values of Kendall's τ , and that the Gumbel copula is only defined for non-negative values of Kendall's τ . Via a rotation, it is however possible to extend these copula families to negative values of τ . An overview on bivariate copula families and their properties, in particular their different tail behavior, can be found e.g. in Brechmann and Schepsmeier (2012).

For sampling, estimation and prediction, we need the joint density/probability mass function of X and Y that is defined as

$$f_{X,Y}(x, y) := \frac{\partial}{\partial x} P(X \leq x, Y = y). \quad (2)$$

In the remainder of this paper, we will refer to f as the joint density function of X and Y .

We now derive formulas for the joint density of X and Y in terms of the copula $C(\cdot, \cdot | \theta)$. We denote by

$$D_1(u, v | \theta) := \frac{\partial}{\partial u} C(u, v | \theta) \quad (3)$$

for $u, v \in]0, 1[$ the partial derivative of the copula with respect to the first variable. Note that this is the conditional density of the random variable

$V := F_Y(Y)$ given $U := F_X(X)$ (Joe, 1997). Table A.6 in Appendix A shows the partial derivative (3) of the Clayton, Gumbel, Frank and Gauss copula (see e.g. Aas et al. (2009) or Schepsmeiner and Stöber (2013) for more details).

Proposition 1 (Density function). *The joint density function $f_{X,Y}$ of a continuous random variable X and a discrete random variable Y is given by*

$$f_{X,Y}(x, y|\theta) = f_X(x) (D_1(F_X(x), F_Y(y)|\theta) - D_1(F_X(x), F_Y(y-1)|\theta)) \quad (4)$$

Proof. By definition

$$\begin{aligned} \frac{\partial}{\partial x} P(X \leq x, Y = y) &= \frac{\partial}{\partial x} P(X \leq x, Y \leq y) - \frac{\partial}{\partial x} P(X \leq x, Y \leq y-1) \\ &= \frac{\partial}{\partial x} C(F_X(x), F_Y(y)|\theta) - \frac{\partial}{\partial x} C(F_X(x), F_Y(y-1)|\theta) \\ &= f_X(x) (D_1(F_X(x), F_Y(y), \theta) - D_1(F_X(x), F_Y(y-1)|\theta)), \end{aligned}$$

which proves the statement. \square

2.2. Marginal distributions

The framework presented in the preceding subsection holds for general pairs of continuous-discrete random variables. In this paper, we focus on joint models for a positive average claim size X and a positive number of claims Y . We model the average claim size X via a Gamma distribution

$$f_X(x|\mu, \delta) = \frac{1}{x\Gamma(\frac{1}{\delta})} \left(\frac{x}{\mu\delta}\right)^{\frac{1}{\delta}} \exp\left(-\frac{x}{\mu\delta}\right) \quad \text{for } x > 0, \quad (5)$$

with mean parameter $\mu > 0$ and dispersion parameter $\delta > 0$. The number of claims Y is a positive count variable, and is modeled as a zero-truncated Poisson (ZTP) distributed variable with parameter $\lambda > 0$,

$$f_Y(y|\lambda) = \frac{\lambda^y}{y!(1-\exp(-\lambda))} \exp(-\lambda) \quad \text{for } y = 1, 2, \dots$$

The generality of our approach easily allows to use other appropriate distributions such as the log-normal for claim severity or the (zero-truncated) Negative Binomial for claim frequency. The models and results presented below can be extended accordingly.

2.3. Joint copula model for average claim sizes and number of claims

Combining the marginal distributions and the copula approach, we obtain the following general model.

Definition 2 (Joint copula model for average claim sizes and number of claims). *The copula-based Gamma and zero-truncated Poisson model for positive average claim sizes X and positive number of claims Y is defined by the joint density function*

$$\begin{aligned}
 f_{X,Y}(x, y|\mu, \delta, \lambda, \theta) \\
 = f_X(x|\mu, \delta) (D_1(F_X(x|\mu, \delta), F_Y(y|\lambda)|\theta) - D_1(F_X(x|\mu, \delta), F_Y(y-1|\lambda)|\theta)) ,
 \end{aligned}
 \tag{6}$$

for $x > 0$ and $y = 1, 2, \dots$

We remark that for discrete marginal distributions, the copula is not unique. However, it is always unique on the range of the marginal distributions. Further, as we consider zero-truncated Poisson variables, the range is large, and the use of copulas does not impose any restrictions.

The model depends on four parameters: the parameters μ, δ (Gamma) and λ (ZTP) for the marginal distributions, and the copula parameter θ . Table 1 displays the parameters and their relationships to the joint distribution.

We now illustrate the influence of the copula parameters and families on the conditional distribution of Y . Therefore, we use the following result.

	average		
	claim size X	number of claims Y	copula family
distribution	Gamma	zero-truncated Poisson	Gauss, Clayton Gumbel, Frank
parameter(s)	$\mu > 0, \delta > 0$	$\lambda > 0$	$\theta \in \Theta$
expectation	$E(X) = \mu$	$E(Y) = \frac{\lambda}{1-e^{-\lambda}}$	—
variance	$Var(X) = \mu^2\delta$	$Var(Y) = \frac{\lambda(1-e^{-\lambda}(\lambda+1))}{(1-e^{-\lambda})^2}$	—

Table 1: Model parameters of the joint distribution for average claim sizes X and number of claims Y . The definition of the copula families is provided in Appendix A.

Proposition 3. *The conditional distribution $Y|X = x$ of the number of claims given an average claim size of x under the copula-based model of Definition 2 is given by*

$$\begin{aligned}
P(Y = y|X = x, \mu, \delta, \lambda, \theta) &= D_1(F_X(x|\mu, \delta), F_Y(y|\lambda)|\theta) \\
&\quad - D_1(F_X(x|\mu, \delta), F_Y(y-1|\lambda)|\theta).
\end{aligned} \tag{7}$$

Proof. This result follows from Proposition 1, as by definition for two random variables X and Y

$$P(Y = y|X = x) = \frac{f_{X,Y}(x, y)}{f_X(x)}.$$

□

Example 4. *We consider a group of policy holders with an expected number of claims of $\lambda = 2.5$. The average claim size is set to $\mu = 1000$ Euro, and we assume that the standard deviation of the average claim size equals 300 Euro, which leads to a dispersion parameter of*

$$\delta = \frac{300^2}{1000^2} = 0.09.$$




Figure 1: Conditional probability mass function of the number of claims Y . Marginal distributions: expected average claimsize $\mu = 1000$ Euro with dispersion parameter $\delta = 0.09$, expected number of claims $\lambda = 2.5$. We condition on an average claim size of $x = 1200$ Euro. Left: Gauss copula with $\tau = 0; 0.1; 0.3; 0.5$. Right: Gauss, Clayton, Gumbel and Frank copula with $\tau = 1/3$.

We condition on an average claim size of $x = 1200$ Euro.

The left panel in Figure 1 displays the conditional probability mass function (7) of $Y|X = x$ for a Gauss copula with four different values of $\tau = 0; 0.1; 0.3; 0.5$. We observe that the four probability mass functions are different, and that for higher values of τ , more mass is assigned to larger values of y . This is due to the dependence of X and Y and the fact that the conditioning value $x = 1200$ Euro is much higher than the expected average claim size of $\mu = 1000$ Euro. The right panel displays the conditional probability mass function for $\tau = 1/3$ and the four different copula families. The choice

of the copula family clearly influences the conditional distribution, although the differences are less pronounced than in the case of varying τ . In summary, this leads to a flexible class of dependence models between a discrete and a continuous variable.

3. Policy loss estimation

Next, we focus on the distribution of the policy loss.

Definition 5 (Policy loss). *For a policy with average claim size X and number of claims Y , the policy loss is defined as the product of the two quantities,*

$$L := X \cdot Y.$$

The policy loss is a positive, continuous random variable, and it depends on the four model parameters displayed in Table 1. A main contribution of this paper is the following result.

Theorem 6. *The distribution of the policy loss L is given by the density function*

$$f_L(l|\mu, \delta, \lambda, \theta) = \sum_{y=1}^{\infty} \left[D_1 \left(F_X \left(\frac{l}{y} | \mu, \delta \right), F_Y(y|\lambda) | \theta \right) - D_1 \left(F_X \left(\frac{l}{y} | \mu, \delta \right), F_Y(y-1|\lambda) | \theta \right) \right] \cdot \frac{1}{y} f_X \left(\frac{l}{y} | \mu, \delta \right)$$

for $l > 0$.

Proof. For simplicity of notation, we omit the model parameters from the formulas. We consider the two-dimensional random variable

$$(L, Y)^\top \in \mathbb{R}^+ \times \{1, 2, \dots\}$$

and derive its joint density mass function. By definition (see Equation (2)),

$$\begin{aligned} f_{L,Y}(l, y) &= \frac{\partial}{\partial l} P(L \leq l, Y = y) \\ &= \frac{\partial}{\partial l} P\left(X \leq \frac{l}{y}, Y = y\right) \end{aligned}$$

as $X = L/Y$. Substituting $x = l/y$, we obtain

$$\begin{aligned} f_{L,Y}(l, y) &= \frac{\partial}{\partial x} P(X \leq x, Y = y) \cdot \frac{\partial x}{\partial l} \\ &= f_{X,Y}\left(\frac{l}{y}, y\right) \cdot \frac{1}{y}. \end{aligned}$$

The result then follows by marginalizing over the discrete random variable Y . □

This implies that we can evaluate the density of the policy loss for all of our four copula models, given a fixed set of parameters. Further, we can evaluate its mean, variance and quantiles based on the density function.

In a first step, we visualize the densities for a given set of parameters in order to investigate the differences between the four copula types and the degree of dependence between the average claim size and the average number of claims. In the simulation study (Section 6) and the case study (Section 7), we show that in the context of regression, the copula-based model leads to a more precise estimation of the policy loss compared to the independence assumption.

We continue Example 4 and use the same parameter settings for the marginal distributions. Figure 2 displays the density of the policy loss for all four copula families and for three different values of Kendall's τ equal to 0.1, 0.3 and 0.5. First, we observe that the distribution is in general left skewed. Further, we observe that the theoretical densities tend to be multimodal, and




Figure 2: Densities of the policy loss for the four copula families and three different values of Kendall’s τ . We use the parameters settings of Example 4.

the multiple modes become more distinct for increasing values of Kendall’s τ . The skewness and multi-modality can be readily explained by Theorem 6. Setting

$$\omega(y, l | \mu, \delta, \lambda, \theta) := \frac{1}{y} P \left(Y = y \mid X = \frac{l}{y}, \mu, \delta, \lambda, \theta \right),$$

the density of the policy loss can be written as an infinite mixture of Gamma distributions

$$f_L(l | \mu, \delta, \lambda, \theta) = \sum_{y=1}^{\infty} \omega(y, l | \mu, \delta, \lambda, \theta) \cdot f_X \left(\frac{l}{y} | \mu, \delta \right).$$

As the individual Gamma densities are skewed, the density of the mixture tends to be skewed, too. Moreover, a mixture of unimodal Gamma densities can be multi-modal as well. The parameter settings of the model influence the number of the modes and how pronounced they are.

Figure 3 displays the expectation μ_L , the 25%-quantile $q_{0.25;L}$ and 75%-quantile $q_{0.75;L}$ of the policy loss as a function of Kendall’s τ . All three quantities are evaluated using numerical integration and numerical root solvers. We use the parameter settings of Example 4 for the marginal distributions.

The solid and dotted lines indicate the mean and the quantiles if we assume that average claim sizes and number of claims are independent. We observe that the independence assumption leads to an overestimation of the policy loss if average claim sizes and number of claims have a negative monotone association (i.e. $\tau < 0$), and it leads to an underestimation if $\tau > 0$. As an example, we compare the expected policy loss under independence (which equals 2723 Euro) to the expected policy loss for $\tau = 0.2$. We obtain 2860 (+5%) Euro (Gauss), 2837 (+4%) Euro (Clayton), 2880 (+6%) Euro (Gumbel) and 2850 (+5%) Euro (Frank).

Based on Figures 2 and 3, we observe a strong dependence of the distribution of the policy loss on the size of Kendall's τ . However, we do not observe a strong dependence on the choice of the copula family.

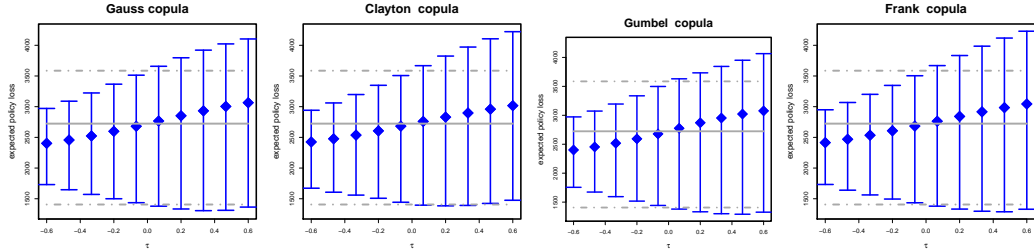


Figure 3: Expected policy loss (blue diamonds) and upper and lower quartiles for the four copula families, seen as a function of Kendall's τ . For negative values of Kendall's τ the Clayton and the Gumbel copula have been rotated. The parameter settings for the marginal distributions are taken from Example 4. The grey solid and dotted lines indicate the expected policy loss and upper/lower quartiles if we assume independence.

4. Copula regression model for average claim sizes and number of claims

In the two previous sections, we modeled average claim sizes and number of claims independently of possible covariates. In order to incorporate covariates, we use the approach by Czado et al. (2012). We extend the joint model (6) for average claim sizes X and number of claims Y by allowing the marginal distributions of X and Y to depend on a set of covariates. More precisely, we apply generalized linear models for the marginal regression problems and combine these with bivariate copula families.

4.1. Model formulation

Let $X_i \in \mathbb{R}_+$, $i = 1, 2, \dots, n$, be independent continuous random variables and let $Y_i \in \mathbb{N}_{>0}$, $i = 1, 2, \dots, n$, be independent discrete random variables. We model X_i in terms of a covariate vector $\mathbf{r}_i \in \mathbb{R}^p$ and Y_i in terms of a covariate vector $\mathbf{s}_i \in \mathbb{R}^q$. The marginal regression models are specified via

$$\begin{aligned} X_i &\sim \text{Gamma}(\mu_i, \delta) && \text{with } \ln(\mu_i) = \mathbf{r}_i^\top \boldsymbol{\alpha}, \\ Y_i &\sim \text{ZTP}(\lambda_i) && \text{with } \ln(\lambda_i) = \ln(e_i) + \mathbf{s}_i^\top \boldsymbol{\beta}. \end{aligned}$$

Here e_i denotes the exposure time. We remark that the covariate vectors \mathbf{r}_i and \mathbf{s}_i can be distinct.

4.2. Parameter estimation

We need to estimate the unknown parameter vector

$$\mathbf{v} := (\boldsymbol{\alpha}^\top, \boldsymbol{\beta}^\top, \theta, \delta)^\top \in \mathbb{R}^{p+q+2} \tag{8}$$

based on n observation pairs (x_i, y_i) . Here, we use maximum-likelihood estimation techniques. By definition, the loglikelihood of the model parameters (8) is

$$\ell(\mathbf{v}|\mathbf{x}, \mathbf{y}) = \sum_{i=1}^n \ln(f_{X,Y}(x_i, y_i|\mathbf{v})) \quad (9)$$

with

$$\mathbf{x} = (x_1, \dots, x_n)^\top \in \mathbb{R}^n \text{ and } \mathbf{y} = (y_1, \dots, y_n)^\top \in \mathbb{R}^n.$$

The maximum likelihood estimates are given by

$$\hat{\mathbf{v}} = \arg \max_{\mathbf{v}} \ell(\mathbf{v}|\mathbf{x}, \mathbf{y}).$$

In general, there is no closed-form solution. Therefore, we have to maximize the loglikelihood numerically. In this paper, we apply the BFGS optimization algorithm (a quasi Newton method) to maximize the loglikelihood (9). As the copula parameter $\theta \in \Theta$ is in general restricted (see Appendix A), we transform θ via a function $g : \Theta \rightarrow \mathbb{R}$ such that $g(\theta)$ is unrestricted. As an example, for the Gauss copula, the copula parameter θ lies in $] - 1, 1[$, and the transformation is defined as

$$g(\theta) = \frac{1}{2} \ln \left(\frac{1 + \theta}{1 - \theta} \right).$$

We then optimize the loglikelihood with respect to $(\boldsymbol{\alpha}^\top, \boldsymbol{\beta}^\top, g(\theta), \delta)^\top$.

Alternatively, we can estimate the model parameters by applying the inference-for-margins (IFM) principle (Joe and Xu, 1996). Here, we proceed in two steps. First, we estimate the marginal regression models for average claim sizes and number of claims via maximum-likelihood. We obtain

estimates

$$\begin{aligned}\widehat{\boldsymbol{\mu}} &= \exp(\mathbf{R}\widehat{\boldsymbol{\alpha}}) \in \mathbb{R}^n \\ \widehat{\boldsymbol{\lambda}} &= \exp(\mathbf{S}\widehat{\boldsymbol{\beta}}) \odot \mathbf{e} \in \mathbb{R}^n\end{aligned}$$

for each observation and an estimate $\widehat{\delta} \in \mathbb{R}$ for the dispersion parameter. Here, $\mathbf{e} \in \mathbb{R}^n$ is the vector of exposure times, and \odot denotes an element-wise multiplication of two vectors. These estimates are used to transform the observations \mathbf{x} and \mathbf{y} to

$$\begin{aligned}u_i &:= F_X(x_i | \widehat{\mu}_i, \widehat{\delta}) \in [0, 1] \\ v_i &:= F_Y(y_i | \widehat{\lambda}_i) \in [0, 1] \\ w_i &:= F_Y(y_i - 1 | \widehat{\lambda}_i) \in [0, 1].\end{aligned}$$

Here, F_X and F_Y are the distribution functions of a Gamma and zero-truncated Poisson variable respectively. In the second step, we optimize the copula parameter θ by maximizing the loglikelihood

$$\widetilde{\ell}(\theta | \mathbf{u}, \mathbf{v}) := \sum_{i=1}^n \ln(D_1(u_i, v_i | \theta) - D_1(u_i, w_i | \theta)).$$

The function $\widetilde{\ell}$ can be maximized numerically. In general, the run-time for the IFM approach is much smaller compared to the maximization of the loglikelihood (9). In initial simulations, the performance of the two methods was very similar. This confirms earlier findings by De Leon and Wu (2011). Therefore, in the simulations study below, we only report the results of the maximum likelihood solution, since it is asymptotically more efficient.

Finally, we note that Czado et al. (2012) recently proposed an extension of the maximization by parts algorithm (Song et al., 2005) to estimate the

regression parameters. These methods could be easily adapted to our model. In this paper, we do not pursue this approach and estimate the parameters via maximum-likelihood.

4.3. Asymptotic distribution of the regression parameters

For the construction of approximate confidence intervals, we use the Fisher information matrix that is defined as

$$\mathcal{I}(\mathbf{v}) := E \left[\frac{\partial \ell(\mathbf{v}|\mathbf{x}, \mathbf{y})}{\partial \mathbf{v}} \cdot \left(\frac{\partial \ell(\mathbf{v}|\mathbf{x}, \mathbf{y})}{\partial \mathbf{v}} \right)^\top \right] \in \mathbb{R}^{(p+q+2) \times (p+q+2)} .$$

Under regularity conditions (see, e.g., Serfling (1980)) one can show that

$$\sqrt{n}(\mathbf{v} - \hat{\mathbf{v}}) \xrightarrow{D} \mathcal{N}_{p+q+2}(\mathbf{0}, \mathcal{I}^{-1}(\mathbf{v})) .$$

Here, \mathcal{N}_k denotes a k -dimensional multivariate normal distribution. For the estimation of the Fisher information, we use the fact that (Lehmann and Casella, 1998)

$$\mathcal{I}(\mathbf{v}) = -E \left[\frac{\partial^2 \ell(\mathbf{v}|\mathbf{x}, \mathbf{y})}{\partial^2 \mathbf{v}} \right] ,$$

and use the observed Fisher information matrix

$$\hat{\mathcal{I}}(\mathbf{v}) := -\frac{\partial^2 \ell(\mathbf{v}|\mathbf{x}, \mathbf{y})}{\partial^2 \mathbf{v}} .$$

This is the Hessian matrix of the loglikelihood function. In our case, it is feasible to compute the second partial derivatives explicitly. Moreover, the BFGS optimization algorithm returns an approximation of the Hessian matrix that is obtained via numerical derivatives. In this paper, we use this approximation to estimate standard errors for the regression coefficients.

4.4. Selection of the copula family

Up to now, the copula class is assumed to be fixed. For the comparison of two copula families, we propose the likelihood-ratio test for non-nested hypotheses by Vuong (1989). This test is appropriate as our models are non-nested, i.e. the regression model for one copula family cannot be obtained via a restriction of the regression model for the other copula family. Let us denote by $\ell^{(1)}, \ell^{(2)} \in \mathbb{R}^n$ the vectors of pointwise loglikelihoods for a model with copula family 1 and 2 respectively. Here, we assume that both models have the same degrees of freedom, i.e. the same number of parameters. We now compute the differences of the pointwise loglikelihood as

$$m_i := \ell_i^{(1)} - \ell_i^{(2)}, \quad i = 1, \dots, n.$$

Denote by

$$\bar{m} = \frac{1}{n} \sum_{i=1}^n m_i$$

the mean of the differences. The test statistic

$$T_V := \frac{\sqrt{n} \cdot \bar{m}}{\sqrt{\sum_{i=1}^n (m_i - \bar{m})^2}}, \quad (10)$$

is asymptotically normally distributed with zero mean and unit variance.

Hence, we prefer copula family 1 to copula family 2 at level α if

$$T_V > \Phi^{-1} \left(1 - \frac{\alpha}{2} \right),$$

where Φ denotes the standard normal distribution function. If

$$T_V < \Phi^{-1} \left(\frac{\alpha}{2} \right),$$

we prefer copula family 2. Otherwise, no decision among the two copula families is possible. We note that it is possible to adjust the test if the two models have different degrees of freedom.

5. Estimation of the total loss

Recall that we model the policy loss $L_i = X_i \cdot Y_i$ for each policy holder via our joint regression model. Its distribution is determined by Theorem 6. In a next step, we are interested in the distribution of the total loss over all policy holders.

Definition 7 (Total loss). *For n policies with average claim sizes X_i and number of claims Y_i (for $i = 1, \dots, n$), the total loss is defined as the sum of the n policy losses*

$$T := \sum_{i=1}^n L_i = \sum_{i=1}^n X_i \cdot Y_i.$$

Just as the individual policy losses, the total loss is a positive, continuous random variable. An application of the central limit theorem leads to the following result.

Proposition 8 (Asymptotic distribution of the total loss). *For n independent policy losses L_1, \dots, L_n with mean μ_{L_i} and variance $\sigma_{L_i}^2$, the asymptotic distribution of the total loss T is normal. For*

$$\sigma_n^2 := \sum_{i=1}^n \sigma_{L_i}^2$$

we have

$$\frac{\sqrt{n}}{\sigma_n} \left(T - \sum_{i=1}^n \mu_{L_i} \right) \xrightarrow{\mathcal{D}} \mathcal{N}(0, 1).$$

For the estimation of the total loss, we need to estimate the means μ_{L_i} and variances $\sigma_{L_i}^2$ of the individual policy losses. This is done by replacing the distribution parameters $\mu_i, \delta, \theta, \lambda_i$ of L_i by their estimates obtained from our joint regression model. Then, the mean and variance can be estimated numerically.

6. Simulation study

We consider a regression problem with $n = 500$ policy groups and the following covariates: age, gender and type of car (A, B or C). We assume that all policy groups contain the same number of persons, which leads to a constant offset. The first column of the design matrices

$$\mathbf{S} := \mathbf{R} := (\mathbf{r}_1, \dots, \mathbf{r}_n)^\top \in \mathbb{R}^{500 \times 5}$$

consists of 1's. This corresponds to marginal regression models with an intercept. The second column corresponds to the covariate age, and is drawn uniformly in the range of 18 and 65. The third column is the dummy variable corresponding to female. Here, the probability of a female policy group is set to 1/2. The last two columns are the two dummy variables corresponding to car type B and car type C. Car type A is represented by the intercept. The probability of a certain car type is set to 1/3. The vector of regression coefficients are defined in Table 2. As an example, in this simulation scenario, a female driver has a negative effect on the average claim size, and a positive effect on the number of claims. We set the constant dispersion parameter δ of the Gamma distribution to $\delta = 0.25$, which implies that the coefficient of

	intercept	age	female	car type B	car type C
average claim size X	-0.50	-0.05	-1.00	+2.00	-0.50
number of claims Y	-1.00	+0.04	+0.30	+0.30	+0.20

Table 2: Regression coefficients for the simulation study.

variation (CV) fulfills

$$\text{CV} = \frac{\sqrt{\text{Var}(X_i)}}{E(X_i)} = \sqrt{\delta} = \frac{1}{2}.$$

We consider the four copula families and three different values of Kendall's τ ,

$$\tau = 0.1; 0.3; 0.5.$$

For each parameter setting, we sample $n = 500$ observations from the true copula regression model, and then fit the regression coefficients and Kendall's τ via maximum likelihood. We consider estimation methods: (1) the independent model: we fit the two marginal regression models and set $\tau = 0$. (2) the joint, copula-based model. We also compute the estimated loss for each of the n policies. We repeat this procedure $R = 50$ times. To evaluate the performance of the two approaches, we consider the following measures for the estimated regression coefficients and the expected policy loss. For a parameter vector $\boldsymbol{\gamma} \in \mathbb{R}^k$ with estimate $\hat{\boldsymbol{\gamma}}$, we are interested in the relative mean squared error which is defined as

$$\text{MSE}_{\text{rel}} := E \left(\frac{1}{k} \sum_{i=1}^k \left(\frac{\gamma_i - \hat{\gamma}_i}{\gamma_i} \right)^2 \right). \quad (11)$$

In the r th iteration step, we obtain an estimate of (11) via

$$\widehat{\text{MSE}}_{\text{rel}}^{(r)} := \frac{1}{k} \sum_{i=1}^k \left(\frac{\gamma_i - \widehat{\gamma}_i^{(r)}}{\gamma_i} \right)^2 .$$

Here, $\widehat{\gamma}^{(r)}$ is the estimate of $\boldsymbol{\gamma}$ obtained in the r th step. In the simulation study, we compare the mean relative mean squared error

$$\overline{\text{MSE}}_{\text{rel}} = \frac{1}{R} \sum_{r=1}^R \widehat{\text{MSE}}_{\text{rel}}^{(r)}$$

computed over all R simulation runs. Note that its variance can be estimated via

$$S_{\overline{\text{MSE}}_{\text{rel}}}^2 = \frac{1}{R} \cdot \frac{1}{R-1} \sum_{r=1}^R \left(\widehat{\text{MSE}}_{\text{rel}}^{(r)} - \overline{\text{MSE}}_{\text{rel}} \right)^2 .$$

Further, we investigate the size of the estimated τ , the estimated total loss, and the value of the Akaike information criterion

$$\text{AIC} := -2\ell(\widehat{\boldsymbol{v}}|\boldsymbol{x}, \boldsymbol{y}) + 2\text{DoF} ,$$

where the Degrees of Freedom (DoF) are the number of estimated parameters in the model. Note that we have $p + q + 2 = 12$ Degrees of Freedom for the joint model and $p + q + 1 = 11$ Degrees of Freedom for the independence model. We prefer the model with the lower AIC score.

Figure 4 displays the results for the Clayton copula. For each quantity that we compute in each of the R simulation runs, we display the mean over all R runs. The means are indicated by a square. The width of each error bar equals twice the standard deviation of the quantity, divided by \sqrt{R} .

The upper row in Figure 4 displays the relative mean-squared error of $\widehat{\boldsymbol{\alpha}}$, $\widehat{\boldsymbol{\beta}}$, and the estimated expected policy loss. Overall, we observe that the

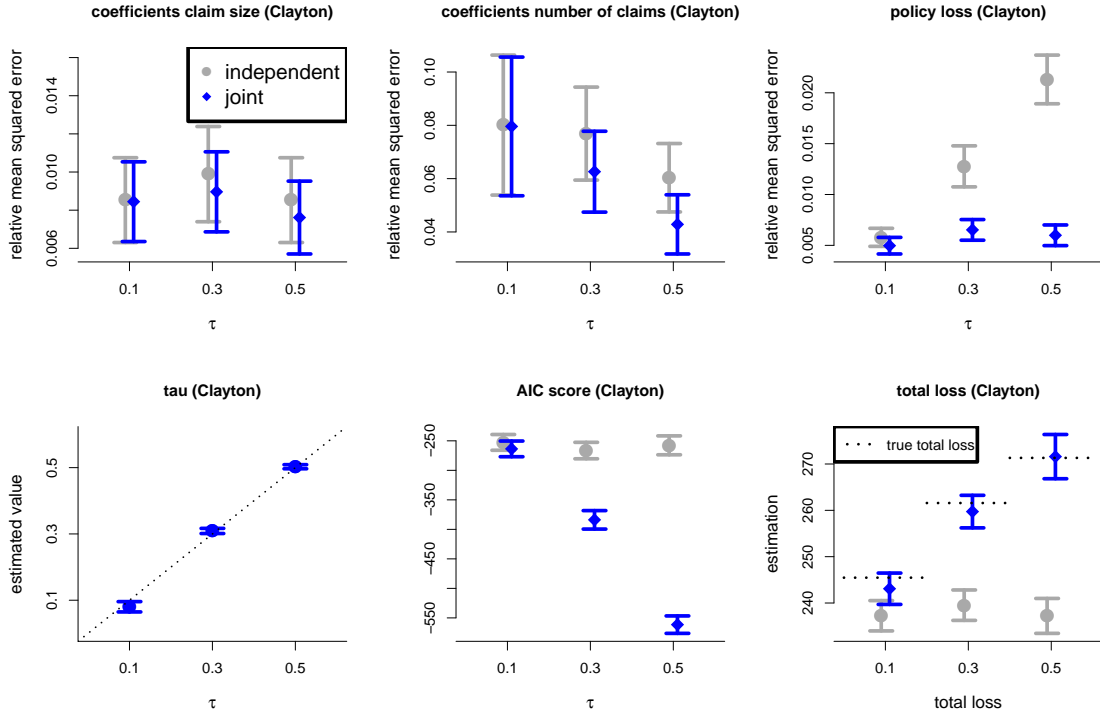


Figure 4: Results of the simulation study for the Clayton copula. Top row: relative mean squared error (11) for the average claim size (left), the number of claims (center) and then policy loss (right). Bottom row: estimated Kendall's τ (left), AIC score (center) and estimated total loss (right). We display the mean over R runs. The width of the whiskers is twice the estimated standard deviation of the mean. Whiskers that are not displayed are too narrow to be visualized.

relative mean squared error for the regression parameters (left and center panel) are not significantly different . For the policy loss however (right panel), the relative mean squared error is lower for the joint, copula-based model, and this improvement becomes more pronounced for higher values of τ .

The first column of the second row displays the estimated value of Kendall's τ . Here, the dashed line indicates the true value of Kendall's τ . We observe that the estimation of τ is very good. Moreover, the AIC score (center panel in the second row) of the joint model is lower than the one of the marginal models . This shows that if joint model is the true model, then we have to use the joint estimation approach, i.e. the dependence cannot be ignored. The right lower panel displays the estimated total loss. The dashed horizontal lines are the true values of the total loss for the respective value of Kendall's τ . We observe that the independence model systematically underestimates the total loss. This confirms the conclusions drawn from Figure 3.

The results for the three other copula families confirm all the findings made for the Clayton copula. We display the results in Appendix B.

7. Case study: car insurance data

We consider data provided by a German insurance company. It contains car insurance data for 7663 German insurance policy groups from the year 2000. It contains seven covariates and information on the exposure time. All seven covariates are categorical. The data was previously analyzed by Czado et al. (2012). Details on the covariates are given in Table 3.

name	description	number of categories
gen	driver's gender	2
rcl	regional class	8
bonus	no-claims bonus	7
ded	type of deductible	5
dist	distance driven	5
age	driver's age	6
const	construction year of the car	7

Table 3: Covariates in the German car insurance data set.

7.1. Marginal models

We first analyze the marginal models. We fit a Gamma regression model for the average claim size, and a zero-truncated Poisson regression model for the number of claims. Next, we investigate the significance of the estimated regression parameters $\hat{\alpha}$ and $\hat{\beta}$. We are interested in those coefficients that are significantly different from 0. Recall (see Section 4.3) that asymptotically, these estimates are normally distributed, and that we can construct approximate confidence intervals using the observed Fisher information. In addition, we adjust the tests for multiple comparisons and the dependence of the estimators (Hothorn et al., 2008). For the number of claims, the covariates age and construction year do not have any significant coefficients on a level of $\alpha = 0.05$. We re-fit the marginal models, leaving out the respective non-significant covariates. Figure 5 displays the joint 95% confidence intervals of the coefficients, showing that the remaining covariates are significant on the 5%-level.




Figure 5: Marginal regression models. Joint 95% confidence intervals for the regression coefficients. Left: Average claim size. Right: number of claims.

7.2. Joint model

We use the covariates selected in Section 7.1 and fit the joint regression model for each copula family. For each pair of copula families, we perform a corresponding Vuong test. Table 4 displays the results. For each pair, we display the copula family that is selected on a $\alpha = 0.05$ -level. In parentheses, we display the value of the Vuong test statistic (10). Note that a value > 2 indicates that we select model 1, and that a value < -2 indicates that we select model 2. We conclude that the Clayton copula is preferred to each of the three other copula families. Therefore, for the remainder of this section, we continue our analysis with the Copula family. The AIC score for the

Clayton model and independence model are

$$\text{AIC}_{\text{clayton}} = 46\,682.35$$

$$\text{AIC}_{\text{independence}} = 46\,921.67.$$

In terms of model comparison, this implies that the copula-based model is more appropriate than the independence model. The estimated value of Kendall's $\tau \pm$ its estimated standard deviation equals

$$\widehat{\tau}_{\text{clayton}} = 0.268 \pm 0.098,$$

which corresponds to a moderate, positive dependence between average claim sizes and number of claims. As a comparison, we note that the estimated value of Kendall's τ for the Gauss copula equals 0.157, which implies that the selection of the copula family has a considerable effect on the estimation of the dependence parameter. Finally, we investigate the impact of this dependence on the estimation of the total loss. For the copula and independence model respectively, we obtain an estimated total loss \pm its estimated

		model 2			
		Gauss	Clayton	Gumbel	Frank
model 1	Gauss	-	<i>Clayton</i> (-10.37)	<i>Gauss</i> (+6.11)	<i>Frank</i> (-5.34)
	Clayton	<i>Clayton</i> (+10.37)	-	<i>Clayton</i> (+9.23)	<i>Clayton</i> (+9.54)
	Gumbel	<i>Gauss</i> (-6.11)	<i>Clayton</i> (-9.23)	-	<i>Frank</i> (-6.54)
	Frank	<i>Frank</i> (+5.34)	<i>Clayton</i> (-9.54)	<i>Frank</i> (+6.54)	-

Table 4: Pairwise Vuong tests. We display the copula family that is selected on a $\alpha = 0.05$ -level. In parentheses, we display the value of the Vuong test statistic (10). A value > 2 indicates that we select model 1, and that a value < -2 indicates that we select model 2.

standard deviation of

$$\begin{aligned}\widehat{E}_{\text{Clayton}}(\mathbb{T}) &= 81\,751.07 \pm 1239.766 \\ \widehat{E}_{\text{independence}}(\mathbb{T}) &= 76\,324.45 \pm 1103.301.\end{aligned}$$

In contrast, the true observed value of the total loss equals 82 256.00. As already illustrated in the simulation study, the negligence of the dependency structure leads to considerably lower estimates of the total loss. In our case, this corresponds to a ratio of

$$\frac{\widehat{E}_{\text{independence}}(\mathbb{T})}{\widehat{E}_{\text{Clayton}}(\mathbb{T})} = 0.934,$$

which indicates a possibly severe underestimation of the independence model in presence of frequency-severity dependence. The more conservative estimate by the copula-based model takes this dependence into account and will thus result in a more appropriate premium rating protecting the insurance company from huge losses in the portfolio.

8. Summary and discussion

In this paper, we model average claim sizes and number of claims if both quantities are dependent. We provide exact distributions of individual policy losses, which tend to be left-skew, and – depending on the parameters of the model – can be multi-modal. Further, we propose a regression approach that models average claim sizes and number of claims in terms of a set of covariates. We showed theoretically and empirically that the explicit incorporation of the dependency in terms of copulae has a profound impact on the estimation of the individual policy loss and the total loss.

Our model explicitly incorporates the discrete structure of the number of claims and allows a flexible class of copula families. This extends previous work that only consider the Gauss copula (Czado et al., 2012; De Leon and Wu, 2011). In our case study, we demonstrated that other copula families are more appropriate.

We stress that our approach does not depend on the particular choice of the marginal distribution or copula family, and it can be extended to other parametric distributions and families (see e.g. Yee and Wild (1996) and Yee (2010) for an overview on appropriate marginal distributions). Moreover, in the case of higher-dimensional mixtures of discrete and continuous random variables, pair-copula constructions (Aas et al., 2009) can be used (see e.g. Panagiotelis et al. (2012) and Stöber et al. (2012)).

In our simulation study, we showed that a model that assumes independence of average claim sizes and number of claims consistently underestimates the total loss of the insurance portfolio implying a severe mispricing of policies. Knowing the true distribution of the policy loss and total loss allows us to correctly assess some risk. This is underpinned in our case study on German car insurance policies. Here, we select relevant covariates for the marginal models and choose the appropriate copula family for the dependence structure using a Vuong test. The data shows a moderate positive dependence. We illustrate that this leads to a more conservative estimation of the total loss, which avoids huge losses in the insurance portfolio and thus possibly filing for bankruptcy. Respecting actuarial prudence therefore requires to take into account possible dependencies between average claim sizes and numbers of claims.

References

- Aas, K., Czado, C., Frigessi, A., Bakken, H., 2009. Pair-copula constructions of multiple dependence. *Insurance: Mathematics and Economics* 44, 182–198.
- Brechmann, E., Schepsmeier, U., 2012. Modeling dependence with C- and D-vine copulas: The R-package CDVine. *Journal of Statistical Software* (to appear) .
- Czado, C., Kastenmeier, R., Brechmann, E., Min, A., 2012. A mixed copula model for insurance claims and claim sizes. *Scandinavian Actuarial Journal* (4) , 278–305.
- De Jong, P., 2009. Modeling Dependence between Loss Triangles Using Copulas. Technical Report. Macquarie University.
- De Leon, A.R., Wu, B., 2011. Copula-based regression models for a bivariate mixed discrete and continuous outcome. *Statistics in Medicine* 30, 175–185.
- Frees, E., Valdez, E., 1998. Understanding relationships using copulas. *North American Actuarial Journal* 2, 1–25.
- Gschlößl, S., Czado, C., 2007. Spatial modelling of claim frequency and claim size in non-life insurance. *Scandinavian Actuarial Journal* 2007, 202–225.
- Haberman, S., Renshaw, A.E., 1996. Generalized linear models and actuarial science. *The Statistician* 45, 407–436.
- Hothorn, T., Bretz, F., Westfall, P., 2008. Simultaneous inference in general parametric models. *Biometrical Journal* 50, 346–363.

- Joe, H., 1997. *Multivariate Models and Dependence Concepts*. Chapman & Hall, London.
- Joe, H., Xu, J., 1996. The estimation method of inference functions for margins for multivariate models. Technical Report, Department of Statistics, University of British Columbia 166.
- Klugman, S., Parsa, R., 1999. Fitting bivariate loss distributions with copulas. *Insurance: Mathematics and Economics* 24, 139–148.
- Krämer, N., Silvestrini, D., 2012. *CopulaRegression: Bivariate Copula Based Regression Models* (version 0.1-3).
- Lehmann, E., Casella, G., 1998. *Theory of Point Estimation*. Springer, New York. 2nd edition.
- Lundberg, F., 1903. *Approximerad framställning af sannolikhetsfunktioner. II. återfrskring af kollektivrisker*. Almqvist & Wiksells Boktr., Uppsala.
- McNeil, A.J., Frey, R., Embrechts, P., 2005. *Quantitative Risk Management: Concepts Techniques and Tools*. Princeton University Press, Princeton.
- Nelsen, R.B., 2006. *An Introduction to Copulas*. Springer, Berlin. 2nd edition.
- Panagiotelis, A., Czado, C., Joe, H., 2012. Pair copula constructions for multivariate discrete data. *Journal of the American Statistical Association* 107, 1063–1072.
- Schepsmeiner, U., Stöber, J., 2013. Derivatives and Fisher information of bivariate copulas. *Statistical Papers* (in press) .

- Serfling, R., 1980. *Approximation Theorems of Mathematical Statistics*. John Wiley & Sons, New York.
- Sklar, A., 1959. Fonction de répartition à n dimensions et leurs marges. *Publ. Inst. Statist. Univ. Paris* 8, 229–231.
- Song, P., 2000. Multivariate dispersion models generated from Gaussian copula. *Scandinavian Journal of Statistics* 27, 305–320.
- Song, P.K., Fan, Y., Kalbfleisch, J., 2005. Maximization by parts in likelihood inference. *Journal of the American Statistical Association* 100, 1145–1166.
- Song, P.X.K., 2007. *Correlated Data Analysis: Modeling, Analytics, and Applications*. volume 365 of *Springer Series in Statistics*. Springer, New York. 1st edition.
- Song, P.X.K., Li, M., Yuan, Y., 2009. Joint regression analysis of correlated data using Gaussian copulas. *Biometrics* 65, 60–68.
- Stöber, J., Czado, C., Hong, H.G., Ghosh, P., 2012. Modeling multiple chronic diseases in the elderly: Longitudinal patterns identified by copula design for mixed responses. Preprint .
- Vuong, Q., 1989. Ratio tests for model selection and non-nested hypotheses. *Econometrica* 57, 307–333.
- Yee, T.W., 2010. The VGAM package for categorical data analysis. *Journal of Statistical Software* 32, 1–34.
- Yee, T.W., Wild, C.J., 1996. Vector generalized additive models. *Journal of the Royal Statistical Society, Series B* 58, 481–493.

Appendix A. Copulae

Table A.5 provides the definition of the four bivariate copula families, their relationship to Kendall's τ and their first partial derivative. Here Φ_2 is the cumulative distribution function of the bivariate standard normal distribution, and Φ is the cumulative distribution function of the univariate standard normal distribution. Further,

$$D_k(x) = \frac{k}{x^k} \int_0^x \frac{t^k}{e^t - 1} dt.$$

denotes the Debye function which is defined for $k \in \mathbb{N}$.

family	copula $C(u, v, \theta)$	range of θ	relationship to Kendall's τ
Gauss	$\Phi_2(\Phi^{-1}(u), \Phi^{-1}(v) \theta)$	$] - 1, 1[$	$\tau = \frac{2}{\pi} \arcsin(\theta) \in \mathbb{R}$
Clayton	$(u^{-\theta} + v^{-\theta} - 1)^{-1/\theta}$	$]0, \infty[$	$\tau = \frac{\theta}{\theta+2} \in]0, \infty[$
Gumbel	$\exp\left(-\left((-\log u)^\theta + (-\log v)^\theta\right)^{1/\theta}\right)$	$]1, \infty[$	$\tau = \frac{\theta-1}{\theta} \in [0, \infty[$
Frank	$-\frac{1}{\theta} \log\left(1 + \frac{(e^{-\theta u} - 1)(e^{-\theta v} - 1)}{e^{-\theta} - 1}\right)$	$\mathbb{R} \setminus \{0\}$	$\tau = 1 - \frac{4}{\theta} [1 - D_1(\theta)] \in \mathbb{R} \setminus \{0\}$

Table A.5: Characteristics of selected copula families.

family	first partial derivative $D_1(u, v \theta)$
Gauss	$\Phi\left(\frac{\Phi^{-1}(v) - \theta \Phi^{-1}(u)}{\sqrt{1-\theta^2}}\right)$
Clayton	$(u^{-\theta} + v^{-\theta} - 1)^{-1/\theta-1} u^{-\theta-1}$
Gumbel	$u^{-1} \exp\left(-\left((-\log u)^\theta + (-\log v)^\theta\right)^{1/\theta}\right)$
Frank	$\frac{e^\theta (e^{\theta v} - 1)}{e^{\theta(u+1)} + e^{\theta(v+1)} - e^\theta - e^{\theta(u+v)}}$

Table A.6: First partial derivative of selected copula families.

Appendix B. Results of the simulation study

We display the results for the Gauss copula (Figure B.6), the Gumbel copula (Figure B.7) and the Frank copula (Figure B.8).

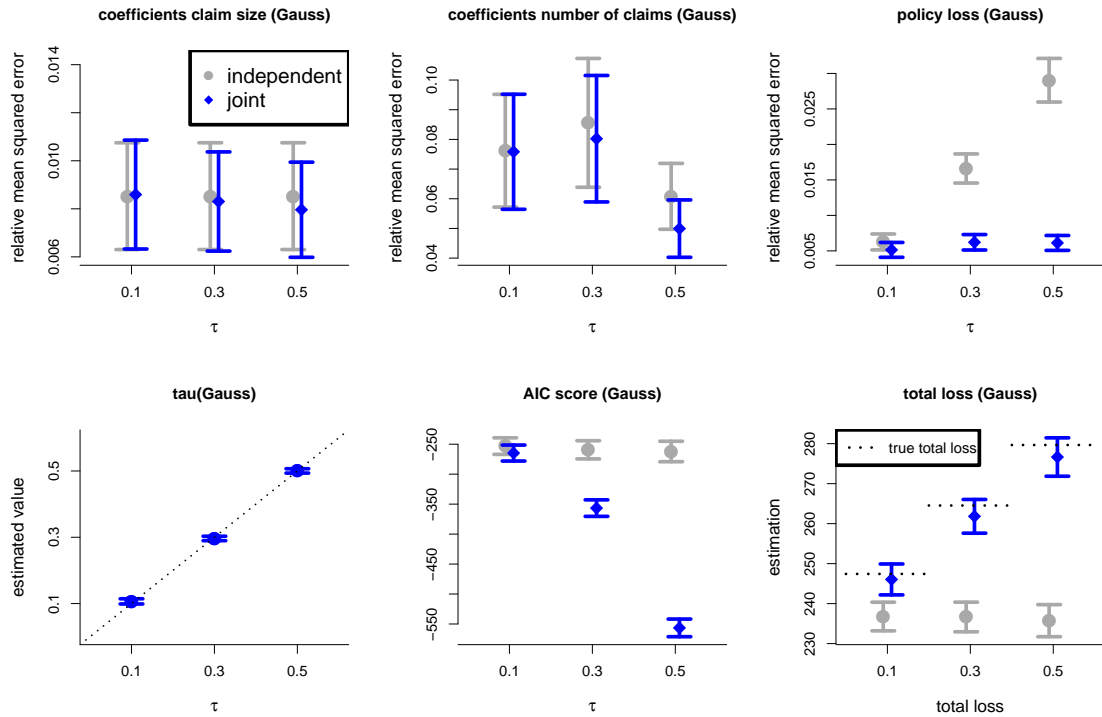


Figure B.6: Results of the simulation study for the Gauss copula. Top row: relative mean squared error (11) for the average claim size (left), the number of claims (center) and then policy loss (right). Bottom row: estimated Kendall's τ (left), AIC score (center) and estimated total loss (right). We display the mean over R runs. The width of the whiskers is twice the estimated standard deviation of the mean. Whiskers that are not displayed are too narrow to be visualized.

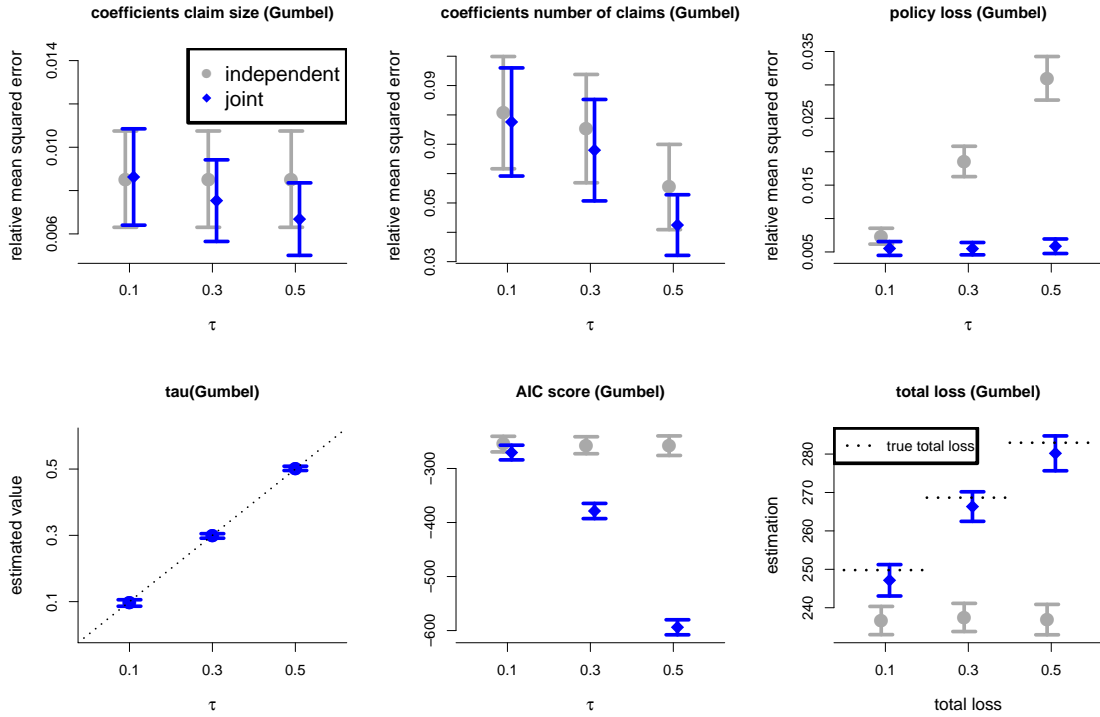


Figure B.7: Results of the simulation study for the Gumbel copula. Top row: relative mean squared error (11) for the average claim size (left), the number of claims (center) and then policy loss (right). Bottom row: estimated Kendall's τ (left), AIC score (center) and estimated total loss (right). We display the mean over R runs. The width of the whiskers is twice the estimated standard deviation of the mean. Whiskers that are not displayed are too narrow to be visualized.

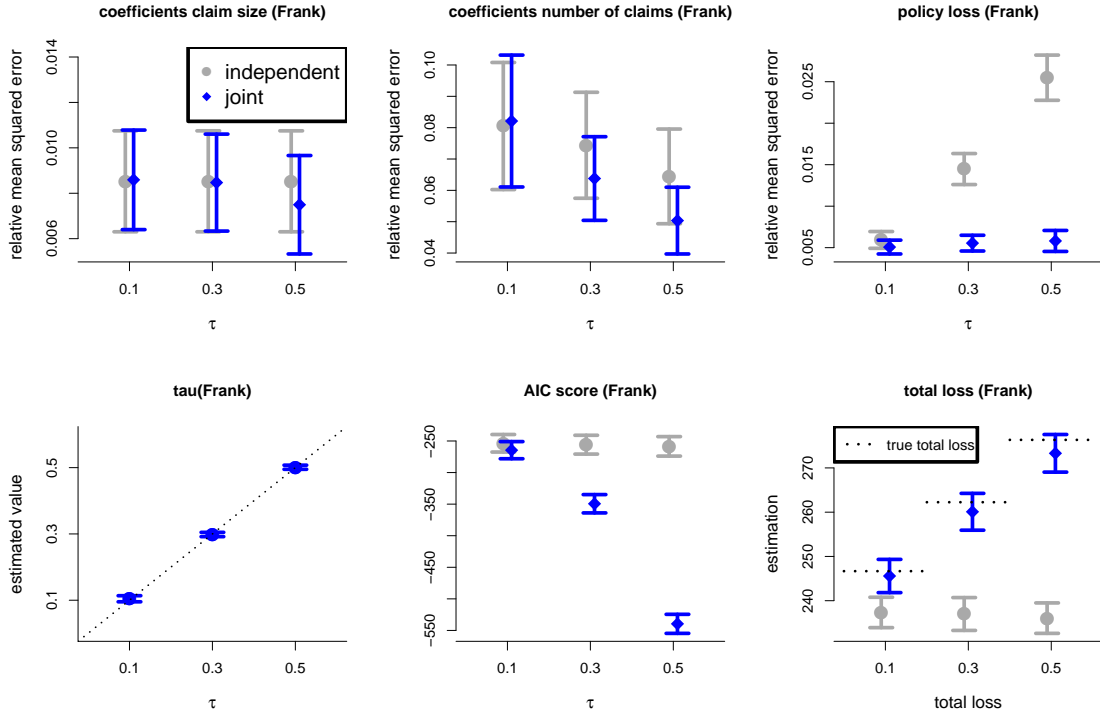


Figure B.8: Results of the simulation study for the Frank copula. Top row: relative mean squared error (11) for the average claim size (left), the number of claims (center) and then policy loss (right). Bottom row: estimated Kendall's τ (left), AIC score (center) and estimated total loss (right). We display the mean over R runs. The width of the whiskers is twice the estimated standard deviation of the mean. Whiskers that are not displayed are too narrow to be visualized.