Lisbon School of Economics & Management
Universidade de Lisboa

MASTER

ACTUARIAL SCIENCE

MASTER´S FINAL WORK

PROJECT

RISK ADJUSTMENT UNDER IFRS 17:

A COPULA-BASED APPROACH TO DEPENDENCE IN NON-LIFE INSURANCE PORTFOLIOS

MARIA MILISA

SUPERVISION:
ALEXANDRA BUGALHO DE MOURA
MARJANA RAJIC

JULY - 2026

Abstract

This research was conducted using anonymized data from a major Croatian insurer, Croatia Osiguranje, where I had the opportunity to analyze the company's non-life insurance portfolio and model the Risk Adjustment (RA) under the new IFRS 17 accounting standard. According to IFRS 17, the RA represents the additional amount an insurer requires to compensate for the uncertainty in future cash flows (CF) arising from non-financial risks. The standard states the RA definition but leaves the actuaries with the freedom to select suitable statistical methods for calculating it.

Insurance portfolios consist of multiple lines of business (LoBs), which may be statistically dependent, meaning that losses across LoBs may be statistically correlated. Because the RA is calculated for the entire portfolio, ignoring the LoB dependence can lead to its significant misestimation. The motivation for this research is to evaluate how modeling dependence between LoBs using copulas can influence RA estimates.

The portfolio I worked on consists of six LoBs, chosen based on their size and tail behavior. To maintain confidentiality, the original CF for each LoB have been scaled and anonymized. Reserves were projected using the Bootstrap Chain Ladder method applied to quarterly paid claims triangles, with estimated future CF adjusted for inflation and discounted using the EIOPA risk-free rates. A marginal distribution was fitted to each LoB and afterwards, copula modeling was applied to capture joint dependence in the portfolio. The RA was calculated following the standard VaR and TVaR approach at an implied confidence level of 75%.

The key empirical finding of this research is that ignoring dependence between LoBs significantly underestimates the aggregated portfolio RA under the same implied confidence level. All copula models considered produced higher results for claim reserves than the models assuming LoB independence. The results are presented as a decision support tool, allowing the company to select the RA model that best reflects their risk appetite and accumulated claims experience.

Keywords: IFRS 17, Risk Adjustment, Loss Reserving, copula, Bootstrap Chain Ladder

Resumo

Este trabalho foi realizado utilizando dados anonimizados de uma das maiores seguradoras croatas, a Croatia Osiguranje, onde tive a oportunidade de analisar a carteira de seguros não vidada companhia e modelar o Risk Adjustment (RA) no âmbito da nova norma contabilísticaIFRS 17. De acordo com a IFRS 17, o RA representa o montante adicional que uma seguradoranecessita para compensar a incerteza nos fluxos de caixa futuros decorrente de riscos não financeiros. A norma define o RA mas deixa aos atuários a liberdade de selecionar os métodosestatísticos adequados para o seu cálculo.

As carteiras de seguros são compostas por múltiplas linhas de negócio (LoBs), que podem ser estatisticamente dependentes, significando que as perdas numa LoB podem estar correlacionadas com as perdas noutra. Uma vez que o RA é calculado para toda a carteira, ignorar adependência entre LoBs pode levar a uma estimativa significativamente incorreta. A motivaçãodeste estudo é avaliar como a modelação da dependência entre LoBs através de cópulaspode influenciar as estimativas do RA.

A carteira analisada é composta por seis LoBs, selecionadas com base na sua dimensão e comportamento da cauda. Para garantir a confidencialidade, os fluxos de caixa originais de cada LoB foram escalados e anonimizados. As reservas foram projetadas utilizando o método Bootstrap Chain Ladder aplicado a triângulos trimestrais de sinistros pagos, com os fluxos decaixa futuros estimados ajustados pela inflação e descontados utilizando as taxas sem risco daEIOPA. Foi ajustada uma distribuição marginal a cada LoB e, posteriormente, foi aplicada amodelação através de cópulas para capturar a dependência conjunta na carteira. O RA foi calculadoseguindo a abordagem padrão de VaR e TVaR a um nível de confiança implícito de 75%.

A principal conclusão empírica desta análise é que ignorar a dependência entre LoBs subestima significativamente o AR agregado da carteira para o mesmo nível de confiançaa implícito.Todos os modelos de cópulas considerados produziram resultados mais elevados para as reservasde sinistros do que os modelos que assumem independência entre LoBs. Os resultados são apresentados como uma ferramenta de apoio à decisão, permitindo à companhia selecionar o modelode RA, nomeadamente o modelo de cópulas a ter em conta na dependência entre LoBs, quemelhor reflita o seu apetite pelo risco e a sua experiência acumulada em sinistros.

Palavras-chave: IFRS 17, Ajustamento de Risco, Reserva de Perdas, cópula, Bootstrap Chain Ladder

Acknowledgments

I would like to express my deepest gratitude to my supervisor, Professor Alexandra Bugalho de Moura. Thank you for your patience and guidance throughout the development of this Thesis.

I would also like to thank my supervisor at Croatia Osiguranje, Marjana Rajic, as well as all of my colleagues, for providing access to the data and the practical insights that greatly helped shape the final outcome of this work.

Finally, to my mom, dad, sister, and partner. Thank you for your unwavering emotional support. It wouldn't have been possible without you.

Contents

List of Abbreviations

AIC	Akaike Information Criterion
BE	Best Estimate
BIC	Bayesian Information Criterion
CL	Chain Ladder
CF	Cash Flows
CoC	Cost of Capital
EIOPA	European Insurance and Occupational Pensions Authority
FCF	Fulfilment Cash Flows
IBNR	Incurred But Not Reported
IFRS	International Financial Reporting Standard
LoB	Line of Business
MLE	Maximum Likelihood Estimation
MOM	Method of Moments
PIT	Probability Integral Transform
PV	Present Value
RA	Risk Adjustment
RBNS	Reported But Not Settled
SCR	Solvency Capital Requirement
SII	Solvency II
TVaR	Tail Value-at-Risk
VaR	Value-at-Risk

Chapter 1

Introduction

This Thesis was written as a Master's Final Work for my master's degree in Actuarial Science at the Lisbon School of Economics and Management. It was developed during my employment at Croatia Osiguranje, a major Croatian insurance company. As the largest and oldest insurer in Croatia, Croatia Osiguranje is the country's market leader in both life and non-life insurance. With a 31.79% market share in non-life insurance, the company manages a large and diverse portfolio with multiple lines of business (LoBs), thus providing an ideal empirical setting for actuarial research.

In this Study, a framework for calculating the IFRS 17 Risk Adjustment (RA) is developed. Under IFRS 17 (IASB, 2017), the RA is defined as the compensation an insurer requires for bearing the uncertainty in the amount and timing of future cash flows (CF) from non-financial risks. In other words, each insurance company owes money for claims it has not yet fully received or paid at any given point in time, and must have that amount stored in its reserve. The true amount of money (future CF) to be paid is uncertain. To better understand what the possible CF amounts could be, a distribution of all possibilities is modeled. The reserve must be sufficient to pay the best estimate (BE) of future CF, which is the mean of the distribution. It also needs to be sufficient to pay the RA, which is the difference between a higher percentile of the future CF distribution and the BE. The RA represents the additional buffer the company holds above the expected future payments. It is meant to cover the uncertainty where actual claims could be higher than expected. According to IASB (2017), this higher percentile of the distribution is referred to as the confidence level. The chosen confidence level in this Study is 75%, which means the total reserve – the BE plus the RA – will be sufficient to cover the actual future CF in 75 out of 100 possible scenarios. In this case, the company accepts a 25% probability that the actual future CF will exceed the reserve.

Although it defines the RA, IFRS 17 does not specify the estimation technique used to determine it (IASB, 2017), thus leaving companies with the freedom to choose their own. This freedom motivated the development of the framework presented in this Study: a rigorous statistical approach that will accurately reflect the risk profile of the company's portfolio.

When analyzing a portfolio with multiple LoBs, a separate distribution of possible future CF for each LoB is not sufficient. The aggregate portfolio distribution is required, which is the distribution of all six LoBs combined. When modeling the combined distribution, dependence between LoBs needs to be accounted for. This is because the LoBs are not independent, but are mutually affected by the same global environment, economy, and legal circumstances. Thus, the portfolio RA is not simply a sum of the individual LoB RAs. To capture the dependence between LoBs, we consider explicit non-linear dependence modeling between LoBs using a copula model.

The central research questions of this Thesis are: "How does the dependence between LoBs affect the RA under IFRS 17, and how sensitive is the RA to the choice of copula family and dependence structure?"

The goal is for the proposed methodology developed in this Thesis to be used as a decision support tool for the company. The results of the Thesis could be used and applied to different data, adjusting for internal experience and level of CF uncertainty by choosing different distributional and copula models.

The data used consists of anonymized quarterly paid claims triangles from Croatia Osiguranje, which tell us the amount of CF paid in each period of time since the occurance of the claim. The triangles cover 80 accident periods (time periods in which the claim occurred) and 80 development periods (time elapsed since the accident) from 2006 Q1 to 2025 Q4 across six LoBs. The methodology of the framework combines four building blocks. First, following England and Verrall (2002), the Bootstrap Chain Ladder (CL) reserving model is applied to the paid claims triangle, simulating 5,000 possible reserve outcomes for each LoB. Second, the EIOPA risk-free rate and forward looking inflation are used to convert the simulated reserve amounts to their present value (PV). Then, a statistical distribution is fitted to each LoB's simulated reserve PV, also referred to as univariate or marginal distribution. Finally, the univariate distributions are aggregated using copulas (International Actuarial Association, 2018), thus capturing the LoB dependence. The RA is then derived from the portfolio distribution produced by the copula model, using both the Value-at-Risk and Tail Value-at-Risk approaches at a 75% implied confidence level (Bulpitt et al., 2017; International Actuarial Association, 2018), as described in detail in Chapter 2.

As well as the many opportunities this topic presents, it also presents several challenges and limitations worth acknowledging. Two biggest challenges in this Study are long projection horizons and LoB dependence. Long projection horizons occur because the company's portfolio is diverse and some LoBs are much longer-tailed than others. The Bootstrap CL methodology fails to simulate them, so certain approximations are required. Dependence between different LoBs presents a challenge when historical data may not fully capture the true joint behavior of losses. These limitations have been addressed within the Study, and the sensitivity of results to the key assumptions is discussed explicitly in Chapter 6.

This Thesis is divided into seven chapters. In Chapter 2, the theoretical framework, literature review and standard calculation methodologies for the RA are introduced. Chapter 3 focuses on the data that was used, namely on the simulation process. Chapters 4 and 5 describe the statistical marginal and joint distributions, respectively, used for describing the company's portfolio. Finally, Chapter 6 focuses on result analysis, and Chapter 7 finishes with the Study's conclusion.

Chapter 2

Motivation and Theoretical Framework

Effective from January 1st, 2023, IFRS 17 was set to establish new principles for the recognition, measurement, presentation, and disclosure of insurance contracts (IASB, 2017). By enhancing comparability at an international level, it brought transparency, quality of financial information, and economic efficiency. Moreover, it reduced the information gap between the providers of capital and the companies receiving the money, as well as helped investors with identifying new opportunities and improving capital allocation globally (Committee on Property and Casualty Insurance Financial Reporting, 2020). It introduced various improvements to the previous standard IFRS 4, while simultaneously creating opportunities for research that can enhance precision in calculating the regulatory requirements. An example of such an opportunity explored in this Study is the RA.

2.1 Fulfilment Cash Flows

The main difference between IFRS17 and the previous standard IFRS4 is that IFRS17 recognizes and measures insurance contracts at a risk-adjusted PV of future CF (IASB, 2017). That is, it introduces a new way of calculating the Fulfilment Cash Flows (FCF). FCF is calculated from three components: the BE (expected PV) of future CF, the discount rate, and the RA. Firstly, the BE of future CF payable by the end of the insurance contract is calculated. Secondly, the discount rate is applied to the BE. Lastly, the topic of this research, the RA, is added to the discounted BE of future CF, in a way that compensates the firm for the non-financial risks existent in the BE (Bulpitt et al., 2017). Simply said, the RA is the additional amount in the company's reserve, added to the PV of the BE of future insurance obligations, to reflect uncertainty. Thus, the FCF can be expressed as:

$$FCF = PV(\mathrm{BE}) + RA. \tag{2.1}$$

2.2 Risk Adjustment Calculation Methods

IFRS 17 does not display explicit requirements to calculate the RA (IASB, 2017), but it is generally calculated in three ways (Bulpitt et al., 2017; International Actuarial Association, 2018; Committee on Property and Casualty Insurance Financial Reporting, 2020): the Value at Risk (VaR) approach, the Tail Value at Risk (TVaR) approach, and the Cost of Capital (CoC) approach. Next, we describe each of these methods.

2.2.1 Value-at-Risk Approach

The VaR approach for calculating the RA is based on the fact that the CF the company will pay out in the future are uncertain. Although the BE of future CF is calculated, the true amount may be higher or lower. The RA is the additional amount of CF that the company could incur, with a chosen probability p. According to IASB (2017), the company selects this probability p, also called the confidence level, based on its risk appetite. The RA is then the difference between the CF amount at quantile p of the distribution and its BE.

According to Klugman et al. (2019) and Embrechts et al. (2002), VaR at probability p is the p-th percentile of a random variable X, or of its corresponding distribution function F_X, denoted by π_p:

$$VaR_p(X) = \pi_p, \tag{2.2}$$

such that

$$\pi_p = \inf\{x \in \mathbb{R} : F_X(x) \geq p\}, \tag{2.3}$$

where $p \in (0, 1)$ represents the implied confidence level under IFRS 17, and X is the discounted PV of future CF.

As mentioned, VaR can be interpreted as the p-th percentile of the loss distribution, such that losses exceeding this amount occur with probability $(1 - p)$. In other words, it is the amount of capital required to cover all losses up to this threshold, with only a $(1 - p)$ probability of losses being larger. Under the VaR approach for calculating the RA, the insurer selects a confidence level, and determines the RA as the excess of VaR over the BE.

$$\mathrm{RA} = \mathrm{VaR}_p(X) - \mathbb{E}(X),$$

where $\mathbb{E}(X)$ denotes the expected PV of future CF or the BE, and p denotes the chosen confidence level.

2.2.2 Tail Value-at-Risk Approach

The biggest limitation of VaR is that it only identifies the loss corresponding to the selected percentile of the portfolio distribution p. It does not provide any information regarding how severe losses beyond this threshold could be. Therefore, extreme losses exceeding the chosen confidence level are not taken into account (Rockafellar and Uryasev, 2002).

TVaR addresses this limitation by incorporating the severity of losses occurring in the tail of the distribution. It is defined for the same percentile p as VaR, but it denotes the expected loss conditional on that loss exceeding the corresponding $\mathrm{VaR}_p(X)$.

Thus, TVaR can be interpreted as the average of all VaR measures above percentile p. Formally, the TVaR of a random variable X at $p \in (0,1)$ is defined as (Klugman et al., 2019):

$$\mathrm{TVaR}_p(X) = \mathbb{E}(X \mid X > \mathrm{VaR}_p(X)), \tag{2.4}$$

where p is the implied confidence level under IFRS 17, and X is the discounted PV of future CF.

In other words, VaR identifies the loss threshold that is exceeded with probability $(1-p)$, while TVaR evaluates the expected severity of losses once the threshold has been exceeded. For this reason, TVaR is always greater than or equal to VaR and is particularly suitable for modeling tail risk in insurance portfolios:

$$TVaR_p(X) > VaR_p(X). \tag{2.5}$$

Finally, the RA using the TVaR approach is calculated as the excess of TVaR over the expected PV of future CF:

$$\mathrm{RA} = \mathrm{TVaR}_p(X) - \mathbb{E}(X),$$

where $\mathbb{E}(X)$ denotes the expected PV of future CF, and p denotes the chosen confidence level.

2.2.3 Cost of Capital Approach

The third and last approach for calculating the RA is the Cost of Capital (CoC) approach. This approach is based on the idea that holding capital as protection from insurance risks is not free. This capital could otherwise be invested elsewhere, or returned to shareholders. The approach requires two inputs: the amount of capital the company needs to hold against its non-financial risks at each future time period, and the rate of return that capital could have earned elsewhere (the CoC rate) (Bulpitt et al., 2017).

The RA is then calculated as the PV of the cost of holding that capital over the lifetime of the insurance contracts:

$$RA = CoC \cdot \sum_{t=0}^{T} \frac{RC_t}{(1+r)^t}, \tag{2.6}$$

where RC_t is the amount of capital the company must hold at time t (risk capital), and CoC is the rate of return that capital could have earned elsewhere (International Actuarial Association, 2018).

This approach is more volatile than the VaR and TVaR approaches because it is sensitive to the choice of discount rate, and the CoC rate (Bulpitt et al., 2017). The CoC rate is typically set to 6% as suggested by Solvency II (European Commission, 2015). However, IFRS 17 does not impose a benchmark for the risk capital, and leaves the final decision up to the company's judgment.

Since this Thesis calculates the RA directly from simulated reserve distributions, the CoC approach is not pursued further. The VaR and TVaR approaches are used instead, as they derive the RA directly from the percentiles of that distribution. They are considered to be a better fit for the distribution-based Bootstrap Chain Ladder (CL) framework developed in the following chapters.

2.3 Literature Review

The methodology developed in this Thesis is based on three different bodies of literature: stochastic loss reserving, copula-based dependence modeling, and IFRS 17 RA calculation. Each of these areas has been studied individually in great extent. The primary contribution of this Study is their aggregation into a single framework applied to a real non-life insurance portfolio. This chapter reviews the key ideas and developments from each area, and identifies the gap in the existing literature.

2.3.1 IFRS 17 RA

IFRS 17 (IASB, 2017) is a new regulatory standard, effective from 2023, that introduces a new framework for insurance contract measurement. It requires the RA to reflect the compensation an insurer requires for bearing non-financial risk uncertainty (IASB, 2017). Unlike its predecessor IFRS 4, IFRS 17 requires a risk-adjusted PV of future CF to be calculated. Bulpitt et al. (2017), International Actuarial Association (2018) and Committee on Property and Casualty Insurance Financial Reporting (2020) provide a practical overview of the three main RA calculation approaches: the VaR approach, the TVaR approach, and the CoC approach. These references cover the theory behind the RA calculation, but do not include a full empirical methodology on real portfolio data.

2.3.2 Loss Reserving

The CL method is one of the most widely used techniques in loss reserving, used for estimating outstanding claim reserves (the total amount still to be paid on claims that have already occurred) of a given portfolio. CL identifies how claims developed over time historically, and projects the same pattern forward to estimate the ultimate claim settlement amount. Mack (1993) established the non-parametric (distribution-free) CL method. This method projects the future reserve estimate and derives its standard error. While this was an important advance, the Mack model still produces only a point estimate of the reserve with an associated standard error: not a full distribution of possible outcomes. For IFRS 17, where a complete distribution is needed to compute the RA, this is insufficient. England and Verrall (2002) later acknowledge this by applying different simulation methods – including the non-parametric bootstrap – to the CL residuals. Instead of just estimating the expected reserve amount, the bootstrap resamples the observed development residuals (how much the actual reserve amount differs from the CL predicted reserve) thousands of times to generate a full predictive distribution of the reserve. The Bootstrap CL is now a standard tool in stochastic reserving and serves as the first building block of the framework developed in this Thesis.

2.3.3 Dependence Modeling

Different LoBs in any insurance portfolio are subject to intersectional dependence. Historically, dependence was modeled with linear correlation matrices. However, it is recognized that linear correlation has fundamental limitations as a dependence measure, as shown by Embrechts et al. (2002). Linear correlation only captures linear dependence, which means that it fails to describe tail dependence. This is often misleading when analyzing insurance portfolios where joint extreme outcomes in the right tail are large drivers of portfolio risk.

As opposed to linear correlation matrices, copulas – introduced by Sklar (1959) – provide a more flexible framework for bivariate distributional modeling by separating the marginal distributions of individual variables from their dependence structure. Nelsen (2006), Joe (2014), Embrechts et al. (2002) and Frees and Valdez (1998) provide an extensive overview of copula theory. Furthermore, International Actuarial Association (2018) discusses copula aggregation in the context of the IFRS 17 RA. It covers the theory behind combining LoBs through copulas – Archimedean and elliptical – and notes that simulation methods such as the Bootstrap and Monte Carlo can be used to generate reserve distributions. However, it does not prescribe or implement any specific methodologies, and leaves that choice to the practitioners.

2.3.4 Research Gap and Thesis Contribution

The existing literature covers loss reserving, non-linear dependence modeling, and IFRS 17 RA calculation as separate topics. The gap acknowledged in this Study is the combination of stochastic reserving with copula-based modeling specifically for the IFRS 17 RA, performed on a non-life insurance portfolio. A specific methodology is developed and applied. The methodology makes a deliberate choice for the RA calculation, uses the non-parametric Bootstrap CL (England and Verrall, 2002), and compares six copula specifications on a portfolio with six LoBs. Moreover, the Study additionally includes EIOPA risk-free rate discounting, forward looking claims inflation, and two separate dependence structures: one based on empirical data and one on Solvency II regulations.

To be more concise, this Thesis combines the Bootstrap CL (England and Verrall, 2002) with EIOPA risk-free rate discounting, forward-looking inflation, univariate distribution fitting, and non-linear dependence modeling. It produces a decision support tool that is made to be applied in practice using the programming language R. The tool quantifies the full range of portfolio RAs computed under different dependence assumptions. As IFRS 17 only came into effect as of January 2023, the standard is still relatively new and published empirical research on specific RA calculation methodologies is limited. Therefore, scientific studies that implement and test concrete methodologies based on real insurance data are useful and needed.

Chapter 3

Data and Reserving Methodology

Croatia Osiguranje has multiple LoBs in its portfolio. Six of them have been chosen for this research based on their size, tail length, and tail heaviness. The six LoBs, ordered by most heavy-tailed to least, are motor third-party liability (`mtpl`), property third party liability (`tpl_prop`), property (`property`), marine, aviation and transport (`mat`), accident (`acc`), and motor (`motor`).

The data for each LoB's CF have been scaled and anonymized as to not represent the actual values of the company's projected future CF. To anonymize the data but keep the original distribution characteristics of the CF and correlations between different LoBs, the original CF have been multiplied by a random variable drawn once from the Lognormal distribution. Since the same scalar was applied to all cells within a given LoB, the shape of the marginal distribution, along with the relative mean, variance, and skewness, is preserved up to an anonymized scale factor.

The data is in the form of a standard actuarial cumulative claims triangle. The actuarial claims triangle is a two-dimensional matrix where the rows represent the accident periods or the dates when the insured events occurred, and the columns represent development lags or the time elapsed since the accident (Mack, 1993). In each cell of the matrix is the cumulative amount paid for the claims that happened in between the given accident period and development lag. The data in this Thesis covers 80 quarterly claim periods (2006 Q1 to 2025 Q4), and 80 quarterly development lags. Table 3.1 shows a simplified example of a cumulative claims triangle with four accident periods and four development lags.

Table 3.1: Illustrative example of a cumulative claims triangle.

Accident period	**2025 Q1**	**2025 Q2**	**2025 Q3**	**2025 Q4**
2025 Q1	$c_{1,1}$	$c_{1,2}$	$c_{1,3}$	$c_{1,4}$
2025 Q2	$c_{2,1}$	$c_{2,2}$	$c_{2,3}$	—
2025 Q3	$c_{3,1}$	$c_{3,2}$	—	—
2025 Q4	$c_{4,1}$	—	—	—

The empty cells in the table represent the unobserved future payments. For example, a claim from 2025 Q4 may still be open, and therefore, its payments are expected to happen in the future. Identifying the amounts in the unobserved cells is the ultimate goal of loss reserving. They are the unknown future CF that the CL method projects (Mack, 1993). The distribution of these future payments is represented by 5,000 simulated reserve scenarios, which will be further explained in Chapter 4.

In non-life insurance, claims are often reported and settled long after the accident happened (International Actuarial Association, 2018). At any given point in time, the insurer owes money for claims it has not yet fully received or paid. This can be due to a serious injury claim involving ongoing medical costs that continue for years, a claim being settled in court for multiple years, a claim being set on hold because of missing documentation etc. They resolve in the claim taking a long time to be resolved and paid for. A LoB with these kinds of claims is said to have a long settlement tail. In this Thesis, the long-tailed LoBs are mtpl, tpl_prop, and property, while the short-tailed are acc, mat, and motor.

The reserve the insurer holds is meant to cover the amounts in the unobserved cells of the triangle. It consists of two parts: the Reported But Not Settled (RBNS), and the Incurred But Not Reported (IBNR). Together, they make the total outstanding claims reserve. Estimating the total outstanding claims reserve is the central challenge of loss reserving, as well as the starting point of this Thesis (England and Verrall, 2002; Mack, 1993). One of the most common methods of estimation is the CL method (England and Verrall, 2002; Mack, 1993). CL uses the historical claim development pattern to project the ultimate claim amount in the future. In its standard form, it produces the reserve estimate as a scalar. However, the IFRS 17 RA requires the full reserve distribution. To do this, we need the residuals, which are the differences between the actual observed claim developments and those predicted by CL. The Bootstrap CL (England and Verrall, 2002) then resamples the observed development residuals with replacement and generates 5,000 alternative reserve scenarios. These scenarios create a statistical distribution of the reserve, as seen in the example illustrated in Figure 3.1.

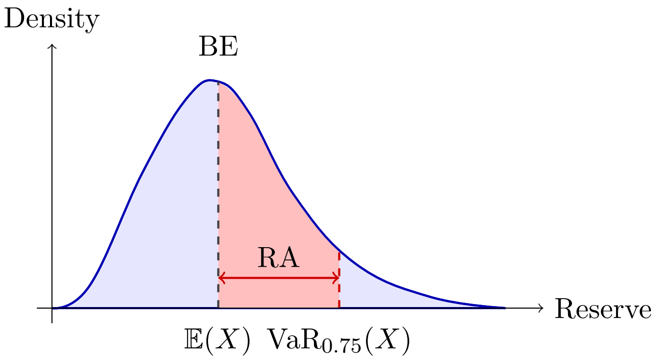

Figure 3.1: Illustrative reserve distribution produced by the Bootstrap CL. The best estimate (BE) is the mean of the distribution $\mathbb{E}(X)$. The RA using thr VaR approach, is the difference between the VaR at the chosen confidence level and the BE.

The non-parametric bootstrap CL (Mack, 1993) was chosen over parametric simulation methods because it does not make any assumptions about the underlying claim development residuals distribution. The final reserve distribution is entirely data-driven and reflects the observed historical variability in claim development. It is also a well-established method in actuarial literature, which makes it a transparent and academically solid choice for a distribution-based RA under IFRS 17 (International Actuarial Association, 2018; Liu and Verrall, 2010; England and Verrall, 2002).

On the other hand, some of the disadvantages of the bootstrap method are that it can be a poor approximation for small samples, it relies heavily on the fact that each sampled variable is independent from one another. Also, when dealing with low frequency and high severity outcomes, the variability of outcomes may not be adequately represented by historical observations (International Actuarial Association, 2018).

In this Thesis, the bootstrap is built on top of the standard Mack CL model (Mack, 1993). The Mack model first estimates *development factors* f_j, which indicate the expected growth of claim amounts from development year j to $j+1$ for a given accident period i:

$$f_j = \frac{\sum_{i=1}^{n-j} C_{i,j+1}}{\sum_{i=1}^{n-j} C_{i,j}},$$

where $C_{i,j}$ denotes the cumulative claims for accident period i at development period j.

Then, it calculates the residuals $\varepsilon_{i,j+1}$, representing the deviations between observed and expected cumulative claim amounts:

$$\varepsilon_{i,j+1} = C_{i,j+1} - f_j C_{i,j},$$

where $C_{i,j}$ denotes the cumulative claim amount for accident quarter i at development lag j. These residuals are subsequently standardised to Pearson scaled residuals before being resampled in the bootstrap procedure (England and Verrall, 2002).

Following England and Verrall (2002), the non-parametric bootstrap method with replacement is applied to the Pearson scaled residuals. Rather than using raw residuals, each residual is scaled by the square root of its fitted value:

$$r_{i,j+1} = \frac{C_{i,j+1} - f_j C_{i,j}}{\sqrt{f_j C_{i,j}}},$$

where $C_{i,j+1}$ is the observed cumulative claims for accident quarter i at development lag $j+1$, $f_j C_{i,j}$ is the CL projection of that cell, and $\sqrt{f_j C_{i,j}}$ is the scaling factor.

Scaling ensures that residuals from different accident periods are comparable. If not applied, larger accident periods would produce larger raw residuals. The scaled residuals are further sampled with replacement, and the simulated triangle is constructed as:

$$C^*_{i,j+1} = f_j C^*_{i,j} + r^*_{i,j+1} \cdot \sqrt{f_j C^*_{i,j}}$$

where $C^*_{i,j+1}$ is the bootstrapped cumulative claim amount for accident quarter i at development lag $j+1$, $C^*_{i,j}$ is the bootstrapped cumulative claim amount at the previous lag, f_j is the CL development factor at lag j, and $r^*_{i,j+1}$ is a Pearson scaled residual sampled with replacement from the observed set $\{r_{i,j+1}\}$. The bootstrap simulation is repeated 5,000 times for all cells in the triangle. Each of the 5,000 simulations produces a complete set of future claims, thus generating 5,000 alternative scenarios of the reserve triangle, which aggregate to the final reserve distribution.

3.1 Calendar Year Decomposition

The Bootstrap CL produces a single scalar outstanding claims reserve $R^{(b)}$ for each simulation $b = 1, \ldots, 5{,}000$. To discount 5,000 reserves from 5,000 bootstrap simulations to their PV, they must be decomposed into CF by calendar year. That is, we need to know how much will be paid in total, and when it will be paid as well. We achieve this by deriving a payment pattern from the observed historical triangle, which tells us what proportion of the total reserve is typically paid in each future period. This pattern is only calculated once from the data, and then applied to all simulated reserves. This way, only the total reserve amount varies across simulations, while the timing of each payment is held fixed at the historical average. A similar simplification was discussed by England and Verrall (2002), who argued that assuming a common payment pattern across accident years when the true patterns are different can result in prediction error for the reserve estimate. However, given that the assumption is applied at a quarterly level in this Thesis, it is considered to be more reasonable. The average of a large number of 80 accident quarters reduces the impact of potential individual quarter's specific pattern. Furthermore, given that CF occurring decades into the future are discounted over long horizons of up to 55 years at the EIOPA risk-free rates, their PV contribution is limited. That is, a cash flow of 1,000 EUR in 2080 discounted at approximately 3.31% per year has a PV of only around 166 EUR today.

To calculate the payment pattern, first, each incremental cell is expressed as a fraction $w_{i,j}$ of its accident quarter's total observed ultimate claims as follows:

$$w_{i,j} = \frac{\Delta C_{i,j}}{\sum_{k=1}^{n} \Delta C_{i,k}},$$

where $\Delta C_{i,j}$ is the incremental payment made in development lag j for accident quarter i, n is the total number of development lags, and the denominator is the sum of all incremental payments for that accident quarter across all development lags. This gives the proportion of total claims paid at each development lag for an accident quarter.

Second, the average fraction at development lag j is computed across accident quarters:

$$\bar{w}_j = \frac{1}{I} \sum_{i=1}^{I} w_{i,j},$$

where I is the total number of accident quarters, 80.

Third, each future unobserved cell (i, j) is mapped to a calendar year $\mathrm{cal}(i, j)$ using the accident quarter index q_i (where $i = 1$ for 2006 Q1 and increases by one each quarter) and the development lag d_j (where $j = 1$ for 2006 Q1 and increases by one each quarter). This procedure makes sure each quarter is correctly mapped to its respective calendar year:

$$\mathrm{cal}(i, j) = 2006 + \left\lfloor \frac{q_i + d_j - 2}{4} \right\rfloor.$$

For example, a cell with an accident quarter equal to q_3 and development lag equal to d_5 would be mapped to the year $\mathrm{cal}(3, 5) = 2007$.

The payment pattern $\tilde{\pi}_y$ is derived using a volume-weighted average across all accident quarters, consistent with the CL methodology. For each development lag j corresponding to calendar year y, the proportion is computed as:

$$\tilde{\pi}_j = \frac{\sum_{i=1}^{I-j} \Delta C_{i,j}}{\sum_{i=1}^{I-j} \sum_j \Delta C_{i,j}},$$

where $\Delta C_{i,j} = C_{i,j} - C_{i,j-1}$ is the incremental paid claims for accident quarter i at development lag j. The denominator sums all incremental payments across the entire observed triangle. The final ratio gives the volume-weighted fraction of total payments occurring at each development lag. By construction, $\sum_j \tilde{\pi}_j = 1$. The CF attributed to calendar year y for bootstrap simulation b is then:

$$CF_y^{(b)} = R^{(b)} \cdot \tilde{\pi}_y,$$

where $R^{(b)}$ is the total reserve from simulation b and $\tilde{\pi}_y$ is the normalized proportion of claims paid in year y.

3.2 Geometric Tail Extension

One limitation of the bootstrap CL model is that there are LoBs whose development is not complete within the 19-year observation window between 2006 and 2025. If we only have 20 years of observations, the bootstrap simulates only 20 future years. This is an issue for long-tailed LoBs, whose payments are expected to extend significantly longer than 20 years. Based on the company's past experience and reserving assumptions, tail extensions are applied to `property` (to 2055), `tpl_prop` (to 2066), and `mtpl` (to 2080). Tail extension is not required for `acc`, `motor`, or `mat`.

The payment pattern is extended beyond the last CL year using geometric tail decay. This approach was chosen because it produces a smooth and monotonically decreasing payment pattern in the distribution tail. This is believed to be a reasonable approximation of the expected behavior of long-tailed claims where the volume of outstanding payments gradually decreases as claims become settled over time.

Let $CF_T^{(b)}$ denote the cash flow in the final CL year T for simulation b, and let T_{end} denote the target end year, which is 2080 for `mtpl`, 2066 for `tpl_prop`, and 2055 for `property` . The number of tail years is then $n = T_{\text{end}} - T$. The tail is constructed in the form of a geometric decay from $CF_T^{(b)}$ to a near-zero target value $\varepsilon > 0$ by year T_{end}. We obtain the per-year decay factor $\delta^{(b)}$ by solving:

$$\delta^{(b)} = \left(\frac{\varepsilon}{CF_T^{(b)}} \right)^{1/n}.$$

We used $\varepsilon = 0.001$, ensuring that tail CF decay to a negligible amount by the final tail year. Since $CF_T^{(b)}$ varies across simulations, the total tail contribution is proportional to the reserve level, thus preserving the distribution properties of the bootstrap simulations.

In practice, tail CF for long-tailed LoBs are not evenly spread. Some years produce significant payments and others produce none. The geometric decay is used as an approximation of the true payments. It produces a smooth tail that avoids the noise of individual year fluctuations. Furthermore, since the tail CF are also subject to large discounting over horizons beyond 50 development years, the PV of the CF being simulated is not significantly impacted.

3.3 Inflation and discount rates

As mentioned before, one of the main differences between IFRS 17 and IFRS 4 is that IFRS 17 requires the future CF to be discounted to their PV (IASB, 2017). In this Study, the EIOPA risk-free rate curve was applied (EIOPA, 2010). The curve covers the period from 2026 to 2175, which includes rates for the analysis of all relevant future CF years.

Let r_t denote the annual risk-free rate for year t from the EIOPA curve. Then, the cumulative discount factor for CF paid at the end of calendar year y is:

$$v(y) = \prod_{t=2026}^{y} \frac{1}{1+r_t},$$

representing the PV at the valuation date (end of 2025 Q4).

In addition, future CF were adjusted for inflation before discounting. Under IFRS 17, future CF are required to incorporate current estimates of future prices, not just the historical levels. Anonymized and rounded projected inflation from the company was used in the following form:

$$i_t = \begin{cases} 6\%, & t = 2026 \\ 5\%, & t = 2027 \\ 4\%, & t = 2028 \\ 2\%, & t \geq 2029 \end{cases},$$

where i_t is the assumed claims inflation rate for calendar year t.

Thus, the cumulative inflation factor for CF paid at the end of year y is:

$$\phi(y) = \prod_{t=2026}^{y} (1+i_t),$$

and the combined net factor applied to each future cash flow is:

$$\psi(y) = \phi(y) \cdot v(y),$$

representing the PV of one unit of a real cash flow paid at the end of year y. For the years 2026, 2027, and 2028, inflation rates are larger than EIOPA discount rates, making the final impact on the PV of future CF positive.

The discounted and inflation-adjusted PV for bootstrap simulation b becomes:

$$PV^{(b)} = \sum_{y \in \mathcal{Y}_{\text{CL}}} CF_y^{(b)} \cdot \psi(y) + \sum_{y \in \mathcal{Y}_{\text{tail}}} CF_y^{(b)} \cdot \psi(y),$$

where $\mathcal{Y}_{\text{CL}}$ is the set of CL calendar years and $\mathcal{Y}_{\text{tail}}$ is the set of tail years. For LoBs without a tail extension, the second sum is zero. This procedure is repeated for each of the 5,000 bootstrap simulations, producing a vector of each of the 5,000 PVs for each LoB:

$$\left\{PV^{(1)},\, PV^{(2)},\, \ldots,\, PV^{(5000)}\right\}.$$

Once aggregated, these PVs form the marginal distribution of the discounted reserve for each LoB.

Chapter 4

Univariate Distribution Modeling

Before aggregating the LoBs into a final portfolio model, a parametric distribution was fitted for each LoB individually. More specifically, a distribution was fit to the 5,000 reserves simulated using the bootstrap method previously described.

4.1 Individual Reserve Distributions

Table 4.1 shows the key statistics of the reserve distribution for each LoB. The mean represents the BE of future CF of that LoB, the VaR represents the value of the portfolio distribution at the 75th percentile, and the net effect represents the combined effect of the discount and inflation rate on the reserve distribution. All monetary values throughout this Thesis are expressed in EUR.

Table 4.1: Discounted and inflation-adjusted reserve distribution by LoB. The Net effect column shows the combined percentage change from applying claims inflation and discount rates to the undiscounted outstanding reserve mean.

LoB	Mean PV	VaR 75%	Net effect
`acc`	2,869,088	3,005,193	+5.14%
`motor`	16,888,653	17,307,544	+4.28%
`mat`	9,178,050	10,350,748	+5.57%
`property`	54,031,682	58,699,587	+4.81%
`tpl_prop`	36,001,051	37,451,949	+5.89%
`mtpl`	278,304,833	287,098,788	+5.39%

Looking at the total oustanding reserve distributions in Table 4.1, `mat` stands out with the highest VaR(75%) relative to its mean, displaying a ratio of a of 1.13 compared to 1.05 for `motor`. Given that marine, aviation, and transport claims are characterized by low frequency and high severity, this is not surprising. Its reserve distribution is skewed to the right with a heavier tail, so the 75th percentile is larger than the mean as compared to other LoBs. On the other hand, despite being the largest LoB by reserve volume, `mtpl` shows a relatively small gap between its mean of 278,304,833 and its VaR(75%) of 287,098,788. This is because the high volume of `mtpl` claims stabilizes the bootstrap distribution in a way that lower-frequency lines (such as `mat`) cannot stabilize.

Furthermore, the net inflation and discount effects reflect the combined impact of applying the forward claim inflation rates (6% in 2026, 5% in 2027, 4% in 2028, and 2% thereafter), and discounting at EIOPA risk-free rates. It is computed as:

$$\text{Net effect} = \left(\frac{\overline{PV}}{\overline{R}} - 1 \right) \times 100\%,$$

where $\overline{PV}$ is the mean discounted and inflation-adjusted PV and $\overline{R}$ is the mean undiscounted total outstanding claims reserve. Once again, because inflation rates from 2026 to 2028 exceed the EIOPA rates for those years, the net effect is positive for all LoBs. In other words, the inflation increase dominates the discounting reduction across the full payment horizon.

4.2 Standalone Risk Adjustment

Before aggregating the portfolio, the standalone RA for each LoB was computed under both the VaR and TVaR approaches at the 75% confidence level:

$$\mathrm{RA}_k^{\mathrm{VaR}} = \mathrm{VaR}_{0.75}(X_k) - \mathbb{E}(X_k),$$

$$\mathrm{RA}_k^{\mathrm{TVaR}} = \mathrm{TVaR}_{0.75}(X_k) - \mathbb{E}(X_k),$$

where k represents the given LoB. The VaR approach is adopted as the primary criterion because it is the most commonly used distribution-based method for calculating the RA in practice (International Actuarial Association, 2018), with TVaR reported as a complementary conservative measure.

Table 4.2: Standalone Risk Adjustment per LoB at VaR(75%) and TVaR(75%).

LoB	Mean PV	VaR(75%)	TVaR(75%)	RA VaR	RA TVaR
acc	2,869,088	3,005,193	3,199,843	155,313	345,880
motor	16,888,653	17,307,544	17,815,620	447,795	955,367
mat	9,178,050	10,350,748	11,955,881	1,072,974	2,770,777
property	54,031,682	58,699,587	64,789,063	4,590,801	10,842,241
tpl_prop	36,001,051	37,451,949	39,228,045	1,502,264	3,296,606
mtpl	278,304,833	287,098,788	296,735,192	9,339,046	18,350,355
Sum	**397,273,357**			**17,108,193**	**36,561,228**

As seen in Table 4.2, due to its largest reserve and longest development tail extending to 2080, `mtpl` yields the highest standalone RA under both approaches. Under the VaR approach, the sum of standalone RAs of 17,108,193 represents the portfolio's perfect positive dependence upper bound. This is the highest possible portfolio RA because all six LoBs simultaneously reach their 75th percentile. Under the TVaR approach, the corresponding sum of 36,561,228 represents the perfect positive dependence upper bound for the average loss in the worst 25% of scenarios. As expected, the TVaR upper bound is significantly higher than the VaR upper bound. We will see in the following chapter that all copula specifications produce a portfolio RA lower than these upper bounds because the true dependence between LoBs is less than 1. This difference between the copula RA and the standalone RA represents the diversification benefit that IFRS 17 paragraph 37 requires to be reflected (IASB, 2017).

4.3 Distribution Fitting

Four candidate distributions were used for fitting a distribution to each LoB: the Gamma, Lognormal, Weibull, and Loglogistic. They were chosen on the basis of two shared properties important for actuarial reserve modeling: right skewness and strict positivity (International Actuarial Association, 2018). Other distributions were considered as well, but these four proved to provide the best fits for the chosen LoBs and the data that was used. The parameters for the Gamma, Weibull, and Loglogistic distribution were estimated by MLE (Maximum Likelihood Estimation), while the parameters for the Lognormal distribution were estimated by MOM (Method of Moments) (Klugman et al., 2019). MLE finds the parameters that maximize the likelihood of observing the given sample, and produces estimators that are consistent and efficient asymptotically. MOM on the other hand, estimates parameters by equating the theoretical moments of the distribution (such as the mean and the variance) to the moments in the sample, and then solves for the parameters of the distribution. For the Lognormal distribution, MOM is equivalent to MLE, which is important to make the information criteria comparable across fitted distributions.

For the Lognormal, the parameters were estimated by MOM on the log scale:

$$\hat{\mu} = \frac{1}{n}\sum_{b=1}^{n} \log PV^{(b)}, \qquad \hat{\sigma} = \sqrt{\frac{1}{n}\sum_{b=1}^{n} (\log PV^{(b)} - \hat{\mu})^2},$$

where $\hat{\mu}$ is the estimated mean of the distribution, and $\hat{\sigma}$ is the estimated standard deviation of the distribution.

The best distribution family was selected by the AIC (Akaike Information Criteria) and BIC (Bayesian Information Criteria). These two information criterion are model selection tools designed for balancing goodness of fit and model complexity (Fabozzi et al., 2014). They are defined as a function of maximized log-likelihood penalized by the number of free parameters. This function rewards better goodness of fit, but at the same time, penalizes unnecessary complexity and potential model overfitting.

AIC is defined as:

$$\text{AIC} = -2\,\ell(\hat{\theta}) + 2k,$$

according to Akaike (1974), where $\ell(\hat{\theta})$ is the maximized log-likelihood evaluated at the MLE parameter estimates $\hat{\theta}$, and k is the number of free parameters in the distribution.

BIC is defined as:

$$\text{BIC} = -2\,\ell(\hat{\theta}) + k/logn,$$

where $\ell(\hat{\theta})$ is the maximized log-likelihood evaluated at the MLE parameter estimates $\hat{\theta}$, and k is the number of free parameters in each candidate distribution (Schwarz, 1978). As displayed in the functions, the log-likelihood holds a negative sign and the penalty holds a positive sign, meaning lower values of AIC and BIC indicate a better fit. Our four distribution families – Gamma, Lognormal, Weibull, and Loglogistic – have the same number of free parameters: $k = 2$. The Gamma is parameterized by shape α and rate β, the Lognormal by μ and σ on the log scale, the Weibull by shape τ and scale θ, and the Loglogistic by shape α and scale β (Klugman et al., 2019). The penalty terms $2k$ in the AIC and $k/logn$ in the BIC, based on the number of free parameters, serve as corrections for overfitting. Without the penalties, a model with more parameters will always achieve a higher log-likelihood and be selected as best, regardless of whether the additional parameters actually improve the fit or not. In our case, the penalties among AIC and among BIC across all distribution families are identical, so the selection is entirely driven by the maximized log-likelihood.

The Loglogistic distribution is selected for all LoBs except for `mtpl`, for which the

Lognormal is selected. This is a reasonable result as it reflects the shape of the bootstrap distribution, which displays moderate right skewness. The distribution reaches a higher peak and displays a heavier tail than that of the Gamma and Weibull distributions, as can be seen in Figure 4.1. However, it is important to note that the margin of selection varies across different LoBs. For `acc` and `motor`, the Loglogistic proves to be a better fit than the Lognormal, by 420.5 and 608.0 BIC points, respectively. For `property`, the Loglogistic is also a better fit than the Lognormal, by 235.3 BIC points. For `mat` and `tpl_prop`, the differences in BIC between the Loglogistic and the Lognormal are relatively small (32.4 and 32.5 points respectively), suggesting that both distributions provide a similar fit. Finally, for `mtpl`, the Lognormal is better than the Loglogistic by only 9 points, indicating that the two distributions are essentially an equivalent fit. We can conclude that the Lognormal and Loglogistic distributions are very close competitors, leaving the ultimate choice to the company, which can contribute its expertise and experience to further improve the purely theoretical selection.

Table 4.3: Distribution fitting results by LoB

LoB	**Distribution**	**AIC**	**BIC**	**Selected**
`acc`	Gamma	139,826.2	139,839.2	
	Lognormal	139,620.2	139,633.3	
	Weibull	142,807.0	142,820.1	
	Loglogistic	139,199.7	139,212.7	✓
`motor`	Gamma	150,239.3	150,252.4	
	Lognormal	150,124.6	150,137.7	
	Weibull	154,735.5	154,748.5	
	Loglogistic	149,516.7	149,529.7	✓
`mat`	Gamma	159,316.7	159,329.7	
	Lognormal	159,232.0	159,245.0	
	Weibull	160,438.8	160,451.8	
	Loglogistic	159,199.6	159,212.6	✓
`property`	Gamma	173,587.9	173,601.0	
	Lognormal	173,450.4	173,463.4	
	Weibull	175,759.6	175,772.6	
	Loglogistic	173,215.0	173,228.1	✓
`tpl_prop`	Gamma	161,747.5	161,760.6	
	Lognormal	161,692.6	161,705.7	
	Weibull	163,329.0	163,342.1	
	Loglogistic	161,660.2	161,673.2	✓
`mtpl`	Gamma	178,867.8	178,880.8	
	Lognormal	178,836.2	178,849.2	✓
	Weibull	180,380.9	180,394.0	
	Loglogistic	178,845.3	178,858.3	

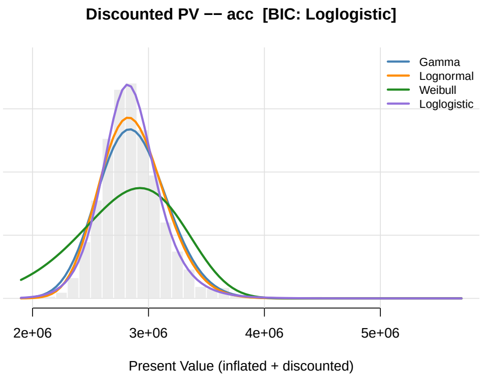

(a) acc

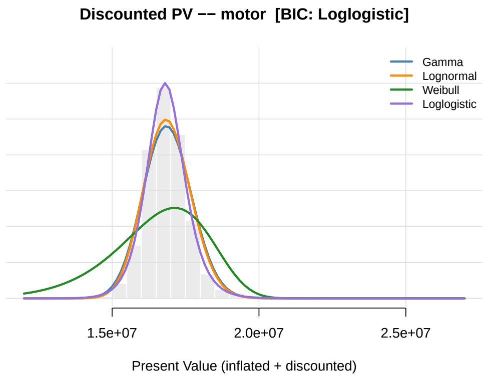

(b) motor

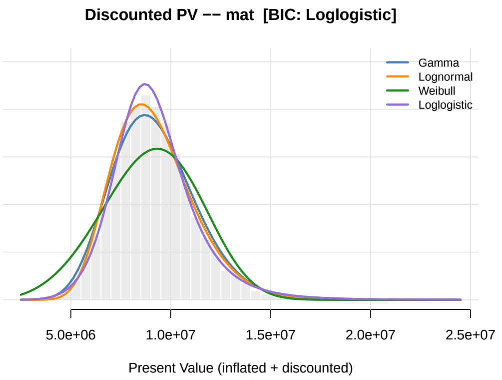

(c) mat

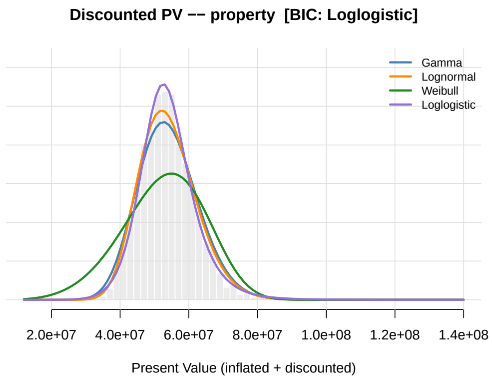

(d) property

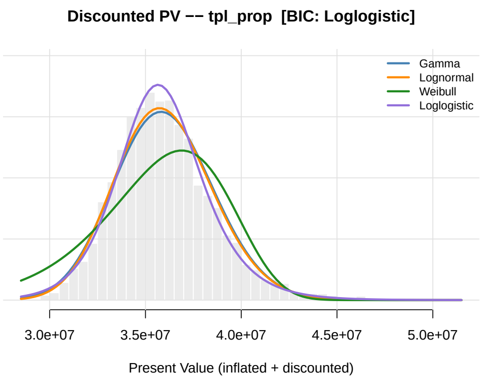

(e) tpl_prop

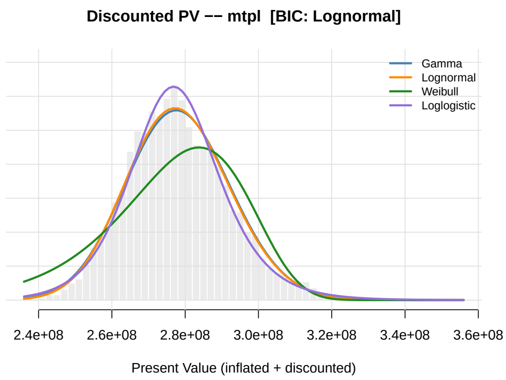

(f) mtpl

Figure 4.1: Fitted univariate distributions for all six LoBs, within the chosen distribution families: Gamma (blue), Lognormal (orange), Weibull (green), Loglogistic (purple).

Chapter 5

Non-linear Dependence Modeling

As the last step in distribution modeling, the univariate distributions fitted in the previous step were aggregated into a copula to account for dependence. Since we will see that the correlations between LoBs in this portfolio are predominantly positive, analyzing the LoBs as if they were independent would lead to an underestimation of portfolio risk and of the RA, which leads us to explicit dependence modeling.

Positive dependence between LoBs can arise due to mutual exposure to inflation, economic conditions, legal environments and catastrophic events. On the other hand, negative dependence can arise in cases when bad outcomes in one LoB coincide with good outcomes in another. For example, one LoB could be sensitive to droughts, while another may be sensitive to freeze damage. Thus, a severe winter that increases freeze claims could coincide with lower drought claims. In such cases, independence would overestimate the portfolio RA. Regardless of whether dependence is positive or negative, accurately modeling it requires a more complicated dependence structure than independence.

Dependence can be modeled through standard linear correlation matrices, but this approach is likely to underestimate non-linear correlation and tail dependence (Embrechts et al., 2002). According to the International Actuarial Association (2018), copulas are an effective tool in addressing these limitations. They are also particularly convenient for modeling non-linear dependence structures. Unlike linear correlation, which only captures the strength of linear dependence between variables, copulas can capture asymmetric dependence, and more importantly, tail dependence. Tail dependence is the tendency for extreme losses across LoBs to occur simultaneously. This is a critical feature for insurance portfolio risk modeling (Nelsen, 2006; Embrechts et al., 2002).

A copula $C : [0,1]^d \to [0,1]$ is a multivariate distribution function with uniform marginal distributions on $[0,1]$. It combines the univariate, marginal distributions into a joint multivariate distribution, while capturing the dependence structure in between variables. In other words, copulas separate the marginal distribution (the behavior of each individual variable), and the dependence structure (the direction in which the variables move together). So, the marginal distribution for each LoB can be chosen independently, while the dependence between LoBs is modeled separately through the copula (Nelsen, 2006; Joe, 2014; Embrechts et al., 2002; Frees and Valdez, 1998).

5.1 Sklar's Theorem

Copulas were first introduced by Sklar (1959). Sklar stated that any multivariate distribution can be decomposed into marginal distributions and a unique copula function describing their dependence (Nelsen, 2006):

Let $X_1, X_2, \ldots, X_d$ be random variables with marginal (univariate) distribution functions $F_1(x_1)$, $F_2(x_2)$, ..., $F_d(x_d)$ and joint distribution function $F(x_1, \ldots, x_d)$. Then there exists a copula $C : [0,1]^d \to [0,1]$ such that for all $(x_1, \ldots, x_d) \in \mathbb{R}^d$:

$$F(x_1, \ldots, x_d) = C\big(F_1(x_1), F_2(x_2), \ldots, F_d(x_d)\big). \tag{5.1}$$

Conversely, if C is a copula and $F_1, \ldots, F_d$ are marginal distribution functions, then the function F defined above is a valid multivariate distribution function with margins $F_1, \ldots, F_d$.

If the marginal distributions are continuous, then the copula C is unique. Otherwise, the copula is uniquely determined on the range of the considered marginal distributions.

5.2 Types of Copulas

Depending on their dependence structure and mathematical properties, copulas can be divided into several copula families. The ones that are widely used in Actuarial Science and risk management are elliptical copulas and Archimedean copulas (Nelsen, 2006; International Actuarial Association, 2018; Joe, 2014).

5.2.1 Elliptical copulas

The elliptical copula family is derived from elliptical distributions, including the Gaussian distribution, which will be applied in this Study. The Gaussian copula is associated with the multivariate normal distribution and is defined as:

$$C_{Gauss}(u_1, \ldots, u_d) = \Phi_P\left(\Phi^{-1}(u_1), \ldots, \Phi^{-1}(u_d)\right), \tag{5.2}$$

where $(u_1, \ldots, u_d) \in [0,1]^d$ are uniform marginal scores, Φ^{-1} denotes the inverse standard univariate normal distribution function and Φ_P denotes the joint multivariate normal distribution with linear correlation matrix P (International Actuarial Association, 2018; Klugman et al., 2019; Nelsen, 2006).

The Gaussian copula is convenient and widely used because no parameter except the linear correlation matrix needs to be estimated. Its one important limitation is that it does not capture tail dependence. Thus, the copula is useful for distributions with low tail dependence, but it can underestimate the probability of extreme joint losses in distributions with higher tail dependence.

5.2.2 Archimedean copulas

Archimedean copulas, defined in Nelsen (2006), Genest and MacKay (1986) and others, are defined by a unique generator function φ:

$$C(u_1, \ldots, u_d) = \varphi^{-1}\left(\varphi(u_1) + \cdots + \varphi(u_d)\right). \tag{5.3}$$

The generator $\varphi : [0,1] \to [0,\infty]$ is a continuous, strictly decreasing function with $\varphi(1) = 0$. The copula is constructed by applying the generator to each uniform score, summing the results, and mapping back through $\varphi^{[-1]}$, the pseudo-inverse of φ, defined as:

$$\varphi^{[-1]}(t) = \begin{cases} \varphi^{-1}(t) & 0 \leq t \leq \varphi(0) \\ 0 & \varphi(0) \leq t \leq \infty. \end{cases}$$

The function C defined in Equation (5.3) is a copula if and only if φ is convex (Nelsen, 2006). When the generator satisfies $\varphi(0) = \infty$, it is called a strict generator, in which case $\varphi^{[-1]} = \varphi^{-1}$ and the copula is called a strict Archimedean copula (Nelsen, 2006). Both the Clayton and Gumbel generators used in this Study are strict, so the pseudo-inverse reduces to the standard inverse.

According to Nelsen (2006), Archimedean copulas have several practical advantages. Due to construction through a single generator function φ, they are simple to compute as well as easy to simulate or extend to any dimension. A large number of copula families belong to their class, and several of them – including the Clayton and the Gumbel – have the interesting property of being defined through only one scalar parameter θ. θ is related to the mean Kendall τ, which is a rank-based correlation measure that quantifies the tendency of two variables to move in the same direction (Kendall, 1938). It is -1 under perfect negative dependence and $+1$ under perfect positive dependence. If the variables are independent, the mean Kendall τ is 0, but the mean Kendall τ being 0 does not necessarily indicate independence.

The Kendall τ takes the mean across all pairs of variables and averages the total variable dependence into one number. The disadvantage here is that all pairs are treated as equally dependent, which is not the case in reality. Although the computation is straightforward and highly applicable, the Archimedean copulas used in this Study ignore the heterogeneity in the pairwise variable correlations. The same limitation exists with

Spearman's ρ, which is also a rank-based dependence measure. It is worth noting that the tail dependence coefficients of the Clayton and Gumbel copulas are also functions of the dependence parameter θ (Nelsen, 2006). Ultimately, tail dependence coefficients could have been estimated directly from the data and used to derive θ. In order to analyze two dependence assumptions – the Empirical and the SII – this approach was not followed in this Study, but could be considered in future research.

The Archimedean copulas used in this paper are the Clayton copula, which captures lower tail dependence, and the Gumbel copula, which captures upper tail dependence.

The Clayton copula is defined as:

$$C_{Clayton}(u, v) = \left(u^{-\theta} + ... + v^{-\theta} - 1\right)^{-1/\theta}, \quad \theta > 0, \tag{5.4}$$

where θ is the dependence parameter (Nelsen, 2006), derived from Kendall's correlation using the following expression:

$$\theta_{\text{Clayton}} = \frac{2\bar{\tau}}{1-\bar{\tau}},$$

where $\bar{\tau}$ is the mean estimated Kendall τ (Nelsen, 2006).

The Gumbel copula is defined as:

$$C_{Gumbel}(u, v) = \exp\left(-\left[(-\ln u)^{\theta} + (-\ln v)^{\theta}\right]^{1/\theta}\right), \quad \theta \geq 1, \tag{5.5}$$

where θ is the dependence parameter (Nelsen, 2006), derived from Kendall's correlation using the following expression:

$$\theta_{\text{Gumbel}} = \frac{1}{1-\bar{\tau}}.$$

where $\bar{\tau}$ is the mean estimated Kendall τ (Nelsen, 2006).

In conclusion, the three copulas – Gaussian, Clayton, and Gumbel – were chosen because each provides a view into a different dependence structure of the portfolio. The Gaussian copula models symmetric dependence and doesn't take tail dependence into account. The Clayton copula models lower tail dependence, and the Gumbel copula models upper tail dependence. The AIC comparison in Section 5.4 will be used to decide which copula family best describes the observed dependence in the data.

5.3 Dependence Parameters

A key decision in copula modeling is the choice of the dependence parameters. In this Thesis, two approaches were used. The first approach is further referred to as the empirical approach because it derives the dependence parameters directly from the observed claims data. The second approach uses the European Commission (2015) QIS5 regulatory correlation matrix to derive the required dependence parameters.

The Gaussian copula requires the full linear Pearson ρ correlation matrix to be used (Nelsen, 2006). In the SII approach, the regulatory SII linear Pearson ρ matrix was used directly. In the empirical approach, the Kendall τ correlation matrix was estimated from the historical data, and was later turned into the linear Pearson ρ correlation matrix.

The Archimedean copulas (Clayton and Gumbel) do not accept the full correlation matrix as input, but require for a single scalar dependence parameter θ to be calculated. In the SII approach, the regulatory linear Pearson ρ correlation matrix was turned into the Kendall τ correlation matrix. From the Kendall τ correlation matrix, the mean Kendall τ was derived. In the empirical approach, the Kendall τ correlation matrix was estimated directly from historical data, which was further used to calculate the mean Kendall τ. The final parameter θ in both approaches was calculated via the inversion formulas in Section 5.2.2.

It is important to note that the three copula families chosen in this Study — Gaussian, Clayton, and Gumbel — represent a small selection from a large number of possible copula families. Other copula families could have been chosen for comparison as well. For example, the Frank copula does not display tail dependence but unlike the Gaussian, belongs to the Archimedean family. Furthermore, the Joe copula displays upper tail dependence, but has a different generator function from the Gumbel copula (Nelsen, 2006; Joe, 2014). The Gaussian, Clayton, and Gumbel were selected because they cover the three fundamentally different dependence structures relevant for non-life insurance portfolios: symmetric dependence in the Gaussian copula, lower tail dependence in the Clayton copula, and upper tail dependence in the Gumbel copula. This choice provides a comprehensive spectrum of dependence assumptions without unnecessarily expanding the number of models.

5.3.1 Empirical Correlation

The empirical Kendall τ correlation matrix was estimated directly from the observed paid claims triangles, and is further referred to as the empirical correlation matrix. The reason why historical triangles were used instead of the bootstrap simulations is that the bootstrap CL method simulates each LoB independently. The matrix that was estimated from the bootstrap simulations produced near-zero Kendall τ across all pairs of LoBs. This indicated that the LoBs were almost independent, which is empirically incorrect.

To estimate the empirical correlation matrix, the final cumulative value to be paid for each given accident quarter once all claims for that quarter are fully developed was estimated. This value is further referred to as the projected ultimate, following Mack (1993). It was computed by applying the remaining CL development factors to the last observed cumulative value. Furthermore, to remove the common upward growth trend that all LoBs share due to inflation and other economic impacts, only quarter-to-quarter percentage changes were used to estimate dependence between LoBs. The Kendall τ correlation matrix was estimated from the 79 detrended observations as:

$$\hat{\tau}_{ij} = \frac{2}{n(n-1)} \sum_{k<l} \operatorname{sgn}\left((x_k - x_l)(y_k - y_l)\right),$$

where $n = 79$, x_k and y_k are the detrended percentage changes for LoBs i and j at quarter k, and $\operatorname{sgn}(\cdot)$ returns $+1$ for concordant pairs, -1 for discordant pairs and 0 for ties (Kendall, 1938).

For the Gaussian copula, each cell in the matrix was converted to the linear Pearson ρ, using the formula:

$$\rho = \sin\left(\frac{\pi}{2}\tau\right),$$

where τ is the mean Kendall τ correlation parameter (Mack, 1993). The resulting empirical and SII correlation matrices are shown in Table 5.1.

For the Archimedean copulas, the mean of all 15 off-diagonal entries $\bar{\tau}$, is computed and used to derive the scalar dependence parameter θ.

Table 5.1: Empirical Kendall τ matrix (top) and linear Pearson ρ matrix converted for Gaussian copulas (bottom).

			Kendall τ			
	`acc`	`motor`	`mat`	`property`	`tpl_prop`	`mtpl`
`acc`	1.000	0.259	0.109	−0.001	0.183	0.439
`motor`	0.259	1.000	0.196	0.266	0.394	0.369
`mat`	0.109	0.196	1.000	−0.048	0.169	0.198
`property`	−0.001	0.266	−0.048	1.000	0.270	0.091
`tpl_prop`	0.183	0.394	0.169	0.270	1.000	0.217
`mtpl`	0.439	0.369	0.198	0.091	0.217	1.000
			Pearson ρ			
`acc`	1.000	0.395	0.171	−0.002	0.284	0.636
`motor`	0.395	1.000	0.304	0.406	0.580	0.548
`mat`	0.171	0.304	1.000	−0.075	0.263	0.306
`property`	−0.002	0.406	−0.075	1.000	0.412	0.143
`tpl_prop`	0.284	0.580	0.263	0.412	1.000	0.335
`mtpl`	0.636	0.548	0.306	0.143	0.335	1.000

The strongest empirical Pearson ρ correlation of 0.636 is between `acc` and `mtpl`. This is intuitive as both LoBs are mainly driven by road traffic claims. They also share exposure to the same legal environment and inflation trends on body injuries. The `motor`–`tpl_prop` pair at 0.580 and `motor`–`mtpl` at 0.548 also show strong positive dependence for similar reasons.

It is important to note that there exist slightly negative correlations between `property` with `acc` and `mat` of (-0.002) and (-0.075). The reason why there is negative correlation between these LoBs is that they are driven by different risk factors. Property claims are influenced by weather events, fire events, construction quality etc., which have little dependence with accident or marine, aviation, and transport claims. Finally, as for the dependence structure of Archimedean copulas, the calculated mean off-diagonal Kendall τ across all 15 pairs of LoBs is 0.2075.

To assess whether the empirical Kendall τ correlations reflect genuine dependence rather than estimation noise, we conducted pairwise significance tests for all 15 pairs of LoBs. For each pair, the null hypothesis of zero dependence was tested against a two-sided alternative:

$$H_0 : \tau_{ij} = 0 \quad \text{vs} \quad H_1 : \tau_{ij} \neq 0.$$

Table 5.2 presents the results.

Table 5.2: Pairwise Kendall τ and p-values of the test to zero correlations. Asterisk denotes significance at the 5% level.

LoB 1	**LoB 2**	**Kendall τ**	**p -value**	**Significant**
`acc`	`motor`	0.259	0.001	$*$
`acc`	`mat`	0.109	0.154	
`acc`	`property`	−0.001	0.990	
`acc`	`tpl_prop`	0.183	0.017	$*$
`acc`	`mtpl`	0.439	0.000	$*$
`motor`	`mat`	0.196	0.010	$*$
`motor`	`property`	0.266	0.001	$*$
`motor`	`tpl_prop`	0.394	0.000	$*$
`motor`	`mtpl`	0.369	0.000	$*$
`mat`	`property`	−0.048	0.534	
`mat`	`tpl_prop`	0.169	0.027	$*$
`mat`	`mtpl`	0.198	0.010	$*$
`property`	`tpl_prop`	0.270	0.000	$*$
`property`	`mtpl`	0.091	0.234	
`tpl_prop`	`mtpl`	0.217	0.005	$*$

Out of the 15 pairs, 11 are statistically significant at the 5% level i.e., the null hypothesis of zero correlations was rejected. The four non-significant pairs — `acc–mat` ($p = 0.154$), `acc–property` ($p = 0.990$), `mat–property` ($p = 0.534$) and `property–mtpl`

($p = 0.234$) — all involve either `property` or `mat`. This is consistent with the empirical near-zero correlations involving these LoBs, as well as their exposure to different risk factors. The remaining 11 pairs are statistically significant, confirming that the dependence structure embedded in the copula model reflects genuine historical correlation between LoBs.

5.3.2 Solvency II Correlation

The second matrix that was used was the linear Pearson ρ matrix, based on the Solvency II QIS5 Technical Specifications EIOPA (2010), SCR.9. Motor family lines (`acc`, `motor`, `mtpl`, `tpl_prop`) are assigned a correlation of 0.50, while all other pairs are assigned the regulatory floor of 0.25. We include `tpl_prop` in the motor family, which is supported by empirical data as supported by the empirical correlation analysis. The data show strong correlation for this LoB with the motor lines, particularly with `motor` at 0.58. The mean off-diagonal SII linear Pearson ρ is 0.350, giving an implied Kendall τ of 0.228, obtained through the following expression (Embrechts et al., 2002):

$$\bar{\tau}_{SII} = \frac{2}{\pi} \arcsin(\bar{\rho}_{SII})$$

The SII linear Pearson ρ correlation matrix is used directly for the Gaussian copula. For the Gumbel and Clayton copula, the calculated implied Kendall τ of 0.228 is used.

The SII matrix is more conservative than the empirical one for the majority of pairs. Exceptions are the correlation between `acc` and `mtpl`, where the empirical correlation of 0.64 is significantly higher than the SII implied one of 0.5, and the correlation between `property` and `acc` and `property` and `mat`, where the empirical correlation is approximately zero, significantly lower than the implied regulatory floor of 0.25.

Figure 5.1: Pearson ρ linear correlation matrices used for the Gaussian copula. Color scale runs from white (zero) to red (positive).

5.4 Comparative Analytics

The choice between the empirical and the SII correlation bases involves a trade-off between statistical precision and regulatory conservatism. The empirical correlation matrix reflects the actual historical dependence in the past 20 years across LoBs. It is an accurate representation of past dependence, but its one disadvantage is that it doesn't necessarily indicate dependence in the future. The correlation structure between LoBs can change, particularly for pairs with near-zero correlations such as `property`–`acc` and `property`–`mat`. For these LoBs, even a small change in the underlying risk environment can significantly alter the true dependence.

Therefore, the regulatory SII matrix was considered as well. The regulatory lower bound for LoB dependence of 0.25 ensures that dependence is not underestimated in scenarios where empirical correlations might understate the true future dependence in between LoBs. By using both matrices, this Study presents the full range of RAs under different copula families, as well as different underlying dependence structures. The final decision is left to the management, who can assess the sensitivity of the RA to the choice of dependence structure.

Before presenting the model quality comparison, it is necessary to define the bootstrap scores used to evaluate the copula fit. Each LoB's 5,000 simulated reserves are transformed into uniform $[0, 1]$ scores by plugging them into their fitted marginal CDF. This is a standard result known as the Probability Integral Transform (Joe, 2014). Intuitively, this converts each simulated value into its percentile rank within the fitted distribution: a value at the 80th percentile of its LoB distribution becomes a score of 0.80. The resulting 5,000 × 6 matrix of uniform scores is then used to evaluate how well each copula captures the joint behavior across all six LoBs. A copula that is a good fit will assign high density to the combinations of scores observed in the bootstrap simulations. Moreover, if the observed scores show that extreme values across LoBs occur simultaneously — that is, many rows of $\mathbf{U}$ have multiple entries close to 1 — then a copula with upper tail dependence such as the Gumbel will fit better than those without.

Table 5.3 shows the comparison of model quality across the considered copulas (Gaussian, Clayton, and Gumbel), as well as the dependence assumptions (empirical and SII). Since all copula parameters are derived from the empirical Kendall τ matrix or the SII regulatory matrix, the number of free parameters k is zero for all six copula specifications. Thus, by the same logic as in 5.2.2, AIC and BIC are identical. For clarity, only the AIC is reported.

The empirical Gumbel copula produces the lowest AIC score of 3,885.9, followed by the SII Gumbel copula at 4,623.6. This result suggests that our LoBs exhibit upper tail dependence. Joint bad outcomes tend to be worse than they would be if modeled with the Gaussian copula family, which does not account for asymptotic dependence in the tail.

Table 5.3 shows the AIC for the six copulas.

Table 5.3: AIC comparison for six copulas

Copula	**Matrix**	**Log-likelihood**	**AIC**
Gumbel	empirical	−1,943.0	**3,885.9**
Gumbel	SII	−2,311.8	4,623.6
Clayton	empirical	−2,490.6	4,981.2
Clayton	SII	−3,087.5	6,175.0
Gaussian	SII	−3,374.4	6,748.9
Gaussian	empirical	−6,404.7	12,809.3

The empirical Gaussian copula performs significantly worse than the others with an AIC of 12,809.3, more than three times higher than the AIC of the Gumbel. This result shows the sensitivity of the Gaussian copula to the slightly negative correlations in the matrix. The slightly negative correlations surrounding the `property` LoB cause the Gaussian copula to assign too low of a density to the observed bootstrap scores. In a Gaussian copula, negative correlation between two LoBs means that one LoB has a high uniform score, and another a low one. This causes the copula to assign too low of a density to scenarios where `property` and other LoBs simultaneously endure high claim amounts. However, the bootstrap simulation shows that `property` and other Lobs occasionally do experience adverse outcomes at the same time. The SII Gaussian copula is slightly less sensitive to this issue because the regulatory floor of 0.25 ensures all correlations are strictly positive. The chosen Archimedean families are less sensitive to the correlation structure in general because they reduce it to the single scalar parameter θ. This is true for the Clayton copula even though, like the Gaussian, it does not display upper tail dependence, as well as the Gumbel.

5.5 Risk Adjustment Analysis

This section presents the portfolio RA computed for all six copula specifications, alongside the independence lower bound and the perfect positive dependence upper bound. The results are reported for both the VaR and TVaR approaches at the 75% confidence level. The confidence level is chosen to be 75% because it is considered to reflect a moderate risk appetite appropriate for a company with a diversified portfolio. Results at the 90%, 95% and 99.5% confidence levels are provided in Appendix C for reference, where 99.5% corresponds to the Solvency II SCR confidence level (European Commission, 2015). The sensitivity of the RA to the choice of copula family and underlying dependence assumptions is discussed. Finally, the implied confidence levels under the independence assumption and the non-linear dependence assumption are compared.

5.5.1 VaR Analysis

Table 5.4 presents the RA under all six chosen copula specifications alongside the independence lower bound and the perfect dependence upper bound.

The independence RA is computed by simulating the independence copula using the same marginal distributions and inverse PIT framework as the six copula specifications. It is a copula that assumes zero dependence between LoBs. The RA using the VaR approach in this case is calculated through the following expression:

$$\mathrm{RA}_{\mathrm{indep}} = \mathrm{VaR}_{0.75}\left(\sum_{k=1}^{6} \tilde{X}_k^{\mathrm{indep}}\right) - \mathbb{E}\left(\sum_{k=1}^{6} \tilde{X}_k^{\mathrm{indep}}\right),$$

where $\tilde{X}_k^{\mathrm{indep}}$ is the simulated discounted PV for LoB k under the independence copula and the portfolio total for each scenario is $\sum_{k=1}^{6} \tilde{X}_k^{\mathrm{indep}}$. This approach ensures that the independence baseline is directly comparable to the six copula specifications, as all seven are computed from the same simulation framework with 100,000 scenarios and identical marginal distributions.

The perfect dependence bound represents the opposite extreme — perfect positive dependence between all LoBs. Under perfect positive dependence, the worst outcomes across all LoBs always occur in the same scenario simultaneously. A key property here is that, under perfect positive dependence, the portfolio VaR equals the sum of the individual LoB VaRs (Nelsen, 2006):

$$\mathrm{VaR}_p(X_1 + \cdots + X_6) = \mathrm{VaR}_p(X_1) + \cdots + \mathrm{VaR}_p(X_6).$$

where X_k is the discounted PV of future CF for LoB k.

All six copula specifications produce a portfolio RA with a value between these two bounds, reflecting a realistic dependence structure that is positive, and neither zero nor one.

The choice of copula can be ultimately left to each company's internal management, and other copulas can be considered as well. The results in Table 5.4 show that all six scenarios produce a higher RA than the independence baseline of 10,891,661, and a lower RA than the perfect positive dependence upper bound of 17,108,193. We can conclude that ignoring cross-LoB dependence would significantly underestimate the true portfolio RA. The RA spans from 12,162,933 under the Empirical Gumbel to 14,284,434 under the SII Clayton, displaying a difference of 2,121,501 that reflects different views on the dependence structure. In this case, the Gumbel copula produces a lower VaR RA at 75% than Clayton because its upper tail dependence effect is concentrated in the extreme tail above the 90th percentile, becoming dominant only at higher confidence levels as shown in Appendix C. Based on their desired risk appetite, the company can choose between these different scenarios. A more risk averse company might gravitate toward the Clayton copulas which produce the highest RAs, while a less risk averse company might gravitate toward the Gumbel copulas, with lower RAs.

In this case, we provide a decision-support tool so that the company can compute the RA for the 6 LoBs, and the type of dependence (copula and correlation structure) chosen.

Table 5.4: Risk Adjustment VaR(75%) under all dependence specifications.

Copula	**Matrix**	**RA VaR(75%)**	**vs Independence**
Independence (lower bound)		10,891,661	—
Gumbel	Empirical	12,162,933	+11.7%
Gumbel	SII	12,454,604	+14.3%
Gaussian	Empirical	13,006,555	+19.4%
Gaussian	SII	13,495,669	+23.9%
Clayton	Empirical	14,023,998	+28.8%
Clayton	SII	14,284,434	+31.2%
Perfect positive dependence (upper bound)		17,108,193	+57.1%

5.5.2 TVaR Analysis

Table 5.5 presents the RA computed using the TVaR approach at the same 75% confidence level. The RA spans from 21,702,609 under the independence assumption to 36,561,228 under the perfect positive dependence, displaying a difference of 14,858,619. It is important to note that among the copula specifications, under the TVaR analysis, the Gumbel copula produces the highest RA of 28,701,172, while under the VaR analysis, it is the Clayton copula at 14,284,434. Thus, apart from the choice of copula and correlation matrix, the choice of the risk measure used for RA computation is significant as well. A company focused only on the 75th percentile might reach different conclusions about the most conservative copula than a company focused on the expected severity of losses beyond the 75th percentile. Both the VaR and TVaR results are valid under IFRS 17.

Table 5.5: Risk Adjustment TVaR(75%) under all dependence specifications.

Copula	**Matrix**	**RA TVaR(75%)**	**vs Independence**
Independence (lower bound)		21,702,609	—
Gumbel	SII	28,701,172	+32.2%
Gumbel	Empirical	28,081,732	+29.4%
Gaussian	SII	27,090,373	+24.8%
Clayton	SII	26,091,661	+20.2%
Gaussian	Empirical	25,830,078	+19.0%
Clayton	Empirical	25,734,472	+18.6%
Perfect positive dependence (upper bound)		36,561,228	+68.5%

5.5.3 Implied Confidence Level

IFRS 17 paragraph 119 requires companies to disclose the confidence level implied by their RA. The implied confidence level p^* is the percentile of the portfolio loss distribution from which the RA is calculated. It can be interpreted as the proportion of possible future scenarios that the total reserve (BE of future CF plus the RA) is sufficient to cover. In our case, a confidence level of 75% means the reserve is sufficient in 75 out of 100 possible scenarios.

Because copulas capture positive dependence between LoBs, the portfolio loss distribution under the copula model has a heavier tail than the independence joint distribution. This happens because adverse outcomes across LoBs tend to occur simultaneously. Therefore, the copula-based RA computed at 75% is equivalent to a higher percentile of the independence distribution, reflecting the additional risk shown by dependence.

Under the VaR approach, the Empirical Gumbel copula (the copula with the lowest

AIC score) RA of 12,162,933 corresponds to the 77.3th percentile of the independence distribution, giving a dependence premium of 2.3 percentage points above the stated 75% level under independence. For the SII Gaussian, the implied confidence level under the independence distribution is 79.5%, resulting in a premium of 4.5 percentage points. This range of 2.3 to 4.5 percentage points reflects the additional risk that dependence modeling captures relative to the independence assumption at the same stated confidence level.

Under the TVaR approach, the Empirical Gumbel copula produces a RA of 28,081,732, corresponding to the 86.8th percentile of the independence TVaR distribution. This gives a dependence premium of 11.8 percentage points above the stated 75% level. For the SII Gaussian, the implied confidence level under the indepencence distribution is 85.3%, resulting in a premium of 10.3 percentage points. This range of 10.3 to 11.8 percentage points is significantly larger than the corresponding VaR premiums of 2.3 to 4.5 percentage points, reflecting the fact that TVaR captures the full severity of outcomes beyond the 75th percentile.

Figure 5.2 shows that while using the data in this Thesis, on the example of the VaR approach, a company that computes its RA under independence at an implied confidence level of 75% is holding less capital in its reserve than a company that accounts for dependence with the same confidence level.

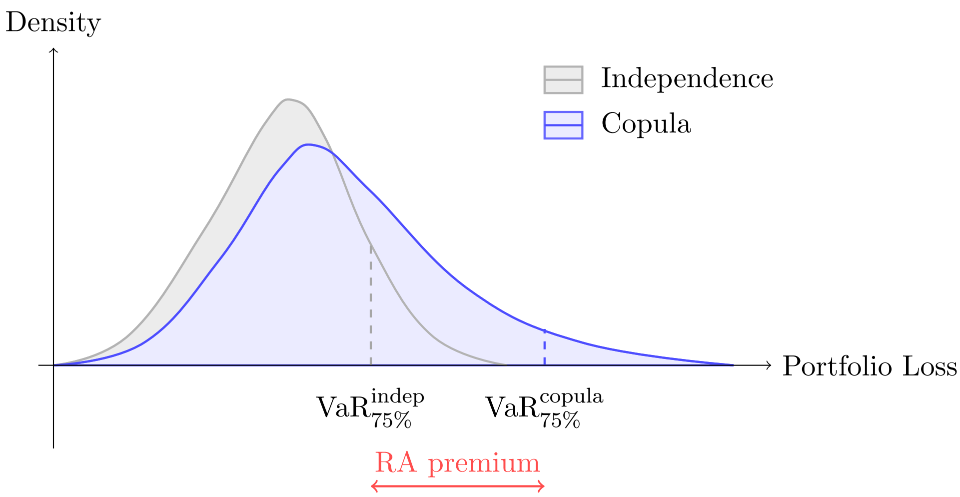

Figure 5.2: Illustrative portfolio loss distributions under independence (grey) and a copula with positive dependence (blue). The copula distribution has a heavier right tail, pushing its 75th percentile further to the right. The gap between the two VaRs of the joint LoB distributions represents the RA premium from dependence modeling.

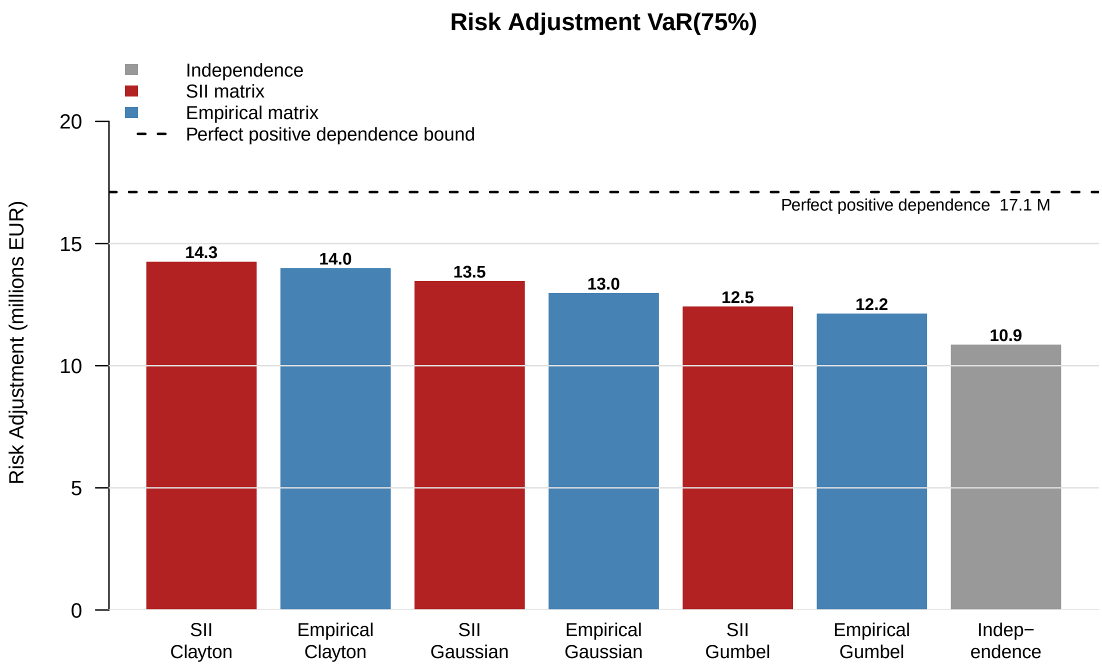

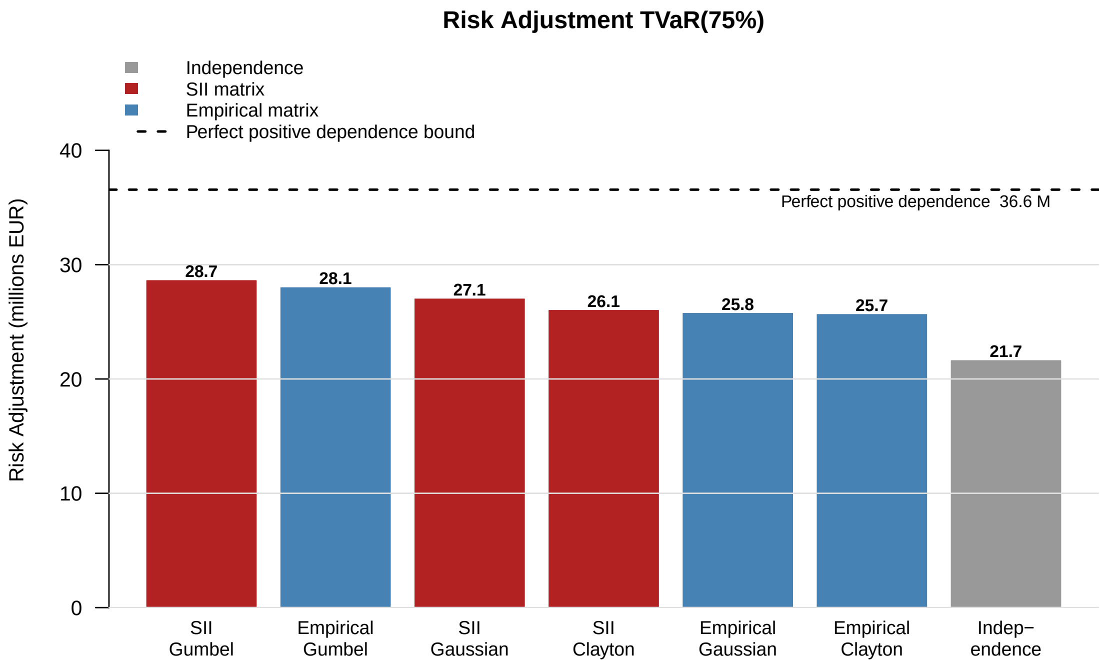

Figure 5.3: RA VaR(75%) (top) and TVaR(75%) (bottom). Red bars correspond to SII-based copulas and blue bars to empirical-based copulas. The dashed horizontal line marks the perfect positive dependence upper bound and the grey bar shows the independence baseline.

Chapter 6

Results and Discussion

This chapter summarizes the findings across the full RA model developed in this Thesis.

The first step of the modeling framework was fitting univariate distributions to the 5,000 simulated bootstrap reserve PVs for each of the six LoBs. Among the univariate models considered, the Loglogistic distribution was selected for five out of six LoBs based on BIC, while the Lognormal was only selected for the longest-tailed LoB, `mtpl` (motor third party liability).

The six LoBs vary significantly in size and tail behavior. Within the scaled and anonymized data used in this Study, `mtpl` has the longest tail. It makes up over 70% of the total reserve and more than half of the standalone RA of 17,108,193. This is expected given that `mtpl` coverage is compulsory for every vehicle in Croatia, claims are frequent, and serious bodily injury cases can result in lifetime annuity payments. The three long-tailed LoBs — `mtpl`, `property` and `tpl_prop` (property third party liability) — account for 89% of the standalone RA, confirming that the RA in this portfolio is largely driven by assumptions on claim settlement time.

The inclusion of claims inflation in the simulated reserve PVs proved to be a crucial modeling distinction. The assumed inflation of 6% in 2026, 5% in 2027, 4% in 2028, and 2% afterwards increased the total portfolio mean PV by approximately 52 million EUR. Even though discounting is supposed to decrease PVs, the high short-term inflation dominated the EIOPA risk-free discount rates, thus producing a net positive effect on the reserve for all LoBs. The effect was largest for `tpl_prop` at +5.89% and smallest for `motor` at +4.28%. The total standalone RA increased from 14,548,238 without inflation to 17,108,193 with inflation, displaying a rise of 17.6%.

The second step included analyzing the dependence parameters needed for the copula specifications considered. Two correlation matrices were used: the empirical Kendall τ matrix, estimated directly from the observed claims data, later converted to a linear Pearson ρ matrix, and the SII linear Pearson ρ matrix. The linear Pearson ρ matrices were used directly as the dependence parameter for the Gaussian copula. For the Clayton and Gumbel copula, which do not accept the linear correlation matrix as the dependence parameter, the mean Kendall τ was derived using the conversion formulas in Section 5.2.2 . The empirical mean Kendall τ correlation analysis confirmed that 11 out of 15 pairs of LoBs show statistically significant positive dependence at the 5% level. The strongest pairs are `acc–mtpl` ($\tau = 0.439$) and `motor–tpl_prop` ($\tau = 0.394$). `Property` stands out with three of its four pairs being statistically insignificant. The SII matrix assigns a correlation of 0.50 to motor family pairs and a regulatory floor of 0.25 to all others.

Six copulas were then fitted using these two matrices across three copula families: Gaussian, Clayton and Gumbel. The AIC comparison favored the Empirical Gumbel copula with an AIC of 3,885.9, suggesting that joint adverse outcomes in the portfolio considered tend to cluster in the upper tail or that adverse years across LoBs tend to happen simultaneously. The Empirical Gaussian copula performed significantly worse than all others with an AIC of 12,809.3. This was because of the slightly negative correlations involving `property` which caused the Gaussian copula to underestimate the probability of joint adverse outcomes. The Gaussian copula predicts that adverse years for `property` and other LoBs do not occur simultaneously, which is in conflict with what the data historically shows.

The key finding of the copula analysis is that ignoring dependence between LoBs in this particular portfolio significantly understates the portfolio RA, calculated using the VaR and TVaR approaches. Under the VaR approach, the independence baseline of 10,891,661 is exceeded by every copula specification, from 11.7% under the Empirical Gumbel to 31.2% under the SII Clayton. Under the TVaR approach, the independence baseline of 21,702,609 is exceeded by every copula specification, from 18.6% under the Empirical Clayton to 32.2% under the SII Gumbel. When adverse years with large claim amounts happen across multiple LoBs at the same time, as indicated by the historical data, the portfolio loss is much worse than what the independence assumption suggests.

The TVaR approach for calculating the RA produced expected results. While the Clayton copulas produce the highest VaR RA at 75%, the Gumbel copulas produce the highest TVaR RA at 75%. The TVAR RA is 28,701,172 under SII Gumbel and 28,081,732 under Empirical Gumbel. This is because TVaR averages all outcomes beyond the 75th percentile. These outcomes are considered worst case scenarios and are most extreme under the Gumbel copula due to its upper tail dependence.

Another interesting finding is that the choice of copula family matters more than the choice of the correlation matrix. Switching between the Gumbel and Clayton families changes the RA by 2,121,501, while switching between the SII and the empirical matrix within the same family changes it by only 260,000 to 490,000. The shape of dependence — whether losses cluster in the upper tail, lower tail or symmetrically — has a larger impact on the RA than the precise pairwise correlations assigned to the LoBs.

These results have a direct practical implication under IFRS 17. Even though all copulas are calibrated at 75%, the dependence they capture pushes the implied confidence level higher. Under the VaR approach, the Empirical Gumbel RA at 75% is equivalent to reserving at 77.3% under independence. The SII Gaussian RA at 75% is equivalent to reserving at 79.5% under independence. A company ignoring dependence would need to add between 2.3 and 4.5 percentage points to its confidence level to match the copula-based VaR RA. Under the TVaR approach, the Empirical Gumbel RA of 28,081,732 corresponds to the 86.8th percentile of the independence distribution. The SII Gaussian TVaR RA at 75% is equivalent to reserving at 85.3% under independence. The TVaR dependence premium ranges from 10.3 to 11.8 percentage points across all six copula specifications.

The results are intentionally presented as a range rather than a single number. From 10,891,661 under independence to 17,108,193 under perfect dependence using the VaR approach, and from 21,702,609 under independence to 36,561,228 under perfect positive dependence using the TVaR approach, the company can see exactly where each copula sits and choose the RA that best reflects its view of the portfolio and its obligations under IFRS 17. It is important to emphasize that IFRS 17 does not prescribe a methodology or a confidence level. Therefore, the model developed in this Study is considered a decision support tool that quantifies the consequences of different assumptions, allowing management to make an informed and documented choice consistent with IFRS 17 requirements.

6.1 Limitations and Future Research

Several limitations of this analysis should be acknowledged.

Most importantly, the available data included 20 years. According to Carrato et al. (2016), stochastic methods are subject to some drawbacks. The main assumption of loss reserving models is that future data are in some way indicated by the past data. However, the claim environment may have changed in a way so that past experience is no longer representative of future experience. Furthermore, the correlation matrix is assumed to be stationary throughout the full observation period from 2006 to 2025, whereas in reality, the dependence between LoBs likely changed due to various economic changes. In future research, with more available data history, it would be interesting to apply a time-varying copula model.

Also, the Archimedean copulas explored in this Study summarize the entire dependence structure with a single scalar number. This number is derived from the mean Kendall τ across all 15 pairs of LoBs, which is a significant simplification and a limitation of using the chosen copulas in practice. As shown in Table 5.1, the pairwise correlations in this portfolio range from near zero for `property–acc` to 0.636 for `acc–mtpl`. Reducing the correlation matrix to a single parameter makes it hard for the copula to distinguish between strongly dependent pairs and almost independent ones. Future research could explore building the model using vine copulas. This model would fit a separate bivariate copula to each pair of LoBs, thus allowing both the copula family and dependence structure to vary across pairs of LoBs. This would represent a more realistic portfolio, but also add significantly more model complexity as well as a larger number of parameters to be estimated from the data.

Finally, due to lack of data, a geometric decay tail extension was applied to CF beyond the observed development period. The geometric tail decay assumes that payments decline smoothly to a negligible amount by the assumed end year. In reality, tail CF for long-tailed LoBs are irregular and do not necessarily follow a geometric pattern. Annuity payments in particular can remain relatively stable for many years before declining. The geometric decay is a pragmatic approximation whose impact on the PV is limited by discounting over long horizons. However, it remains an assumption that introduces model uncertainty into the results. Possible future research could explore alternative tail specifications, such as a flat tail assuming constant payments beyond the CL horizon, a company-specific tail factor derived from historical data, or a stochastic tail model that captures the uncertainty in both the amount and timing of distant CF. Comparing the sensitivity of the RA to different tail assumptions would provide valuable additional insight for practitioners.

Chapter 7

Conclusion

The results from this Study successfully answered the research questions. Accounting for dependence between LoBs when calculating the IFRS 17 RA matters significantly.

When LoBs are treated as independent, adverse years in one LoB can be offset by good years in another. This produces a narrow portfolio loss distribution for the data considered. Under independence, the VaR RA is 10,891,661 and the TVaR RA is 21,702,609. When positive dependence is introduced through a copula, adverse years can happen simultaneously across all LoBs, making the distribution wider and its tail heavier. This pushes the 75th percentile further from the mean and produces a larger RA. Under the AIC-winning Empirical Gumbel copula, these rise to 12,162,933 under the VaR approach and 28,081,732 under the TVaR approach. Under the same implied confidence level of 75%, the RAs accounting for dependence between LoBs is higher. Therefore, a company that would ignore dependence in this data, might be holding a reserve that sits at a lower effective percentile of the true portfolio distribution than it seems. It would effectively be holding a smaller amount than necessary in its reserve, and have less protection against adverse outcomes than its stated confidence level suggests.

The choice of copula can further be updated according to the company's internal expectations about the current and/or future claims circumstances. If there is an ongoing event – such as a major storm, a period of high inflation, or a change in legal environment affecting body injury claims – that is likely to simultaneously increase claims across multiple LoBs, the Gumbel copula might be the appropriate choice. On the other hand, if there are assumptions about a single LoB experiencing the worst outcomes, the Clayton copula might be the appropriate choice. In this case, even the independence assumption can be considered. The Gaussian copula can serve as a middle ground, if normal conditions are expected, with moderate dependence across LoBs. Finally, the perfect dependence case can be considered as a worst case scenario or a stress test that displays a theoretical upper limit for the RA. It is a useful reference point, even if it is never reported as the final RA.

To conclude, the implementation of this methodology in R constitutes a tool for the company to compute the RA under different model assumptions, including dependence between LoBs. The company user can easily include new copula models, as well as univariate models to the computational tool, as well as consider any number of Lobs as necessary.

Bibliography

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723.

Bulpitt, T., Gilchrist, A., Kay, A., and Simpson, P. (2017). IFRS 17 — risk adjustment. Technical report, Milliman.

Carrato, A., McGuire, G., and Scarth, R. (2016). A practitioner's introduction to stochastic reserving. Technical report, Institute and Faculty of Actuaries.

Committee on Property and Casualty Insurance Financial Reporting (2020). IFRS 17 risk adjustment for non-financial risk for property and casualty insurance contracts. Draft Educational Note 220063, Canadian Institute of Actuaries.

EIOPA (2010). QIS5 technical specifications. Technical report, European Insurance and Occupational Pensions Authority, Frankfurt.

Embrechts, P., McNeil, A. J., and Straumann, D. (2002). Correlation and dependence in risk management: Properties and pitfalls. *Risk Management: Value at Risk and Beyond*, pages 176–223.

England, P. D. and Verrall, R. J. (2002). Stochastic claims reserving in general insurance. *British Actuarial Journal*, 8(3):443–518.

European Commission (2015). Commission delegated regulation (EU) 2015/35: Article 39 — cost-of-capital rate. Official Journal of the European Union. Supplementing Directive 2009/138/EC (Solvency II).

Fabozzi, F. J., Focardi, S. M., Rachev, S. T., and Arshanapalli, B. G. (2014). *The Basics of Financial Econometrics: Tools, Concepts, and Asset Management Applications.* John Wiley & Sons, Hoboken, New Jersey.

Frees, E. W. and Valdez, E. A. (1998). Understanding relationships using copulas. *North American Actuarial Journal*, 2(1):1–25.

Genest, C. and MacKay, R. J. (1986). The joy of copulas: Bivariate distributions with uniform marginals. *The American Statistician*, 40(4):280–283.

Gesmann, M., Murphy, D., Zhang, Y. W., Carrato, A., Wuthrich, M., Concina, F., and Dal Moro, E. (2026). *ChainLadder: Statistical Methods and Models for Claims Reserving in General Insurance*. R package version 0.2.21.

IASB (2017). IFRS 17 insurance contracts. International Accounting Standards Board, London.

International Actuarial Association (2018). *Risk Adjustments for Insurance Contracts under IFRS 17*. International Actuarial Association, Ottawa.

Joe, H. (2014). *Dependence Modeling with Copulas*. CRC Press, Boca Raton.

Kendall, M. G. (1938). A new measure of rank correlation. *Biometrika*, 30(1/2):81–93.

Klugman, S. A., Panjer, H. H., and Willmot, G. E. (2019). *Loss Models: From Data to Decisions*. Wiley, Hoboken, New Jersey, 4 edition.

Liu, H. and Verrall, R. J. (2010). Bootstrap estimation of the predictive distributions of reserves using paid and incurred claims. *Variance*, 4(2):121–135.

Mack, T. (1993). Distribution-free calculation of the standard error of chain ladder reserve estimates. *ASTIN Bulletin*, 23(2):213–225.

Nelsen, R. B. (2006). *An Introduction to Copulas*. Springer, New York, 2 edition.

Rockafellar, R. T. and Uryasev, S. (2002). Conditional value-at-risk for general loss distributions. *Journal of Banking and Finance*, 26(7):1443–1471.

Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464.

Sklar, A. (1959). Fonctions de répartition à n dimensions et leurs marges. *Publications de l'Institut de Statistique de l'Université de Paris*, 8:229–231.

Appendix A: R Code

This appendix presents the key R code blocks used in the Study. All code was developed in R version 4.5.0.

A.1 Bootstrap Chain Ladder

The Bootstrap CL is applied to each LoB using the `ChainLadder` package. The `process.distr = "gamma"` argument specifies the process variance distribution used to simulate the over-dispersion in the residuals (Gesmann et al., 2026).

```
boot <- BootChainLadder(tri_fixed,
                        R    = 5000,
                        process.distr = "gamma")
reserve_sim <- as.numeric(boot$IBNR.Totals)
```

A.2 Inflation and Discount Factors

The cumulative inflation factor and EIOPA discount factor are pre-computed for each future calendar year and stored in lookup vectors.

```
compute_infl <- function(cal_year) {
  years_seq <- seq(BASE_YEAR, cal_year)
  i_seq <- sapply(as.character(years_seq), function(y) {
    v <- INFL_RATES[y]
    if (!is.na(v)) as.numeric(v) else INFL_LONGRUN
  })
  prod(1 + i_seq)
}

compute_df <- function(cal_year) {
  years_seq <- seq(BASE_YEAR, cal_year)
  r_seq <- sapply(as.character(years_seq), function(y) {
```

```
    v <- rates[y]
    if (!is.na(v)) as.numeric(v) else last_eiopa_rate
  })
  1 / prod(1 + r_seq)
}
```

The combined net factor applied to each cash flow is:

```
net_factor <- infl_lookup[y] * df_lookup[y]
```

A.3 Payment Pattern and Present Value

The payment pattern is derived from the observed incremental triangle and applied to each bootstrap simulation to decompose the total reserve into calendar year CF, which are then discounted and inflated.

```
pay_pat   <- get_payment_pattern(tri_fixed)
cl_years  <- as.integer(names(pay_pat))

for (s in seq_len(N_SIMS)) {
  cf_cl     <- reserve_sim[s] * pay_pat
  pv_cl     <- discount_cf(cf_cl)
  seed      <- as.numeric(cf_cl[as.character(max_cl_yr)])
  cf_tail   <- build_tail_cf(seed, tail_cal_yrs)
  pv_sim[s] <- pv_cl + discount_cf(cf_tail)
}
```

A.4 Marginal Distribution Fitting

All four distributions are fitted by MLE.

```
# Gamma MLE
fit_g   <- MASS::fitdistr(x, "gamma")
alpha_g <- fit_g$estimate["shape"]
beta_g  <- fit_g$estimate["rate"]

# Lognormal (MOM = MLE)
mu_l    <- mean(log(x))
sigma_l <- sd(log(x))
```

```
# Weibull MLE
fit_w   <- MASS::fitdistr(x, "weibull")
shape_w <- fit_w$estimate["shape"]
scale_w <- fit_w$estimate["scale"]

# Loglogistic MLE on normalised data
x_med <- median(x)
ll_loglogis_fn <- function(par) {
  alpha <- par[1];  beta <- par[2]
  if (alpha <= 0 || beta <= 0) return(Inf)
  x_n  <- x / x_med
  vals <- log(alpha) - log(beta) +
          (alpha - 1) * log(x_n / beta) -
          2 * log(1 + (x_n / beta)^alpha)
  if (any(!is.finite(vals))) return(Inf)
  -sum(vals)
}
fit_ll <- optim(par    = c(alpha = 5, beta = 1),
                fn     = ll_loglogis_fn,
                method = "L-BFGS-B",
                lower  = c(1e-4, 1e-4),
                control = list(maxit = 2000))

# Back-transform scale to original units
sc_ll <- fit_ll$par["beta"] * x_med
sh_ll <- fit_ll$par["alpha"]
```

The Loglogistic distribution is fitted on normalized data $x/\mathrm{median}(x)$ for numerical stability. The scale parameter is transformed back to the original units after convergence.

A.5 Copula Construction

Six copulas are constructed from two correlation matrices and three copula families. The Archimedean theta parameters are derived from the mean Kendall τ using closed-form inversion formulas.

```
# Archimedean theta from mean Kendall tau
clay_theta_emp <- 2 * tau_emp_mean / (1 - tau_emp_mean)
gumb_theta_emp <- 1 / (1 - tau_emp_mean)
```

```
# Six copula objects
cop_sii_gauss <- normalCopula(sii_param, d, "un")
cop_sii_clay  <- claytonCopula(clay_theta_sii, d)
cop_sii_gumb  <- gumbelCopula(gumb_theta_sii, d)
cop_emp_gauss <- normalCopula(emp_param, d, "un")
cop_emp_clay  <- claytonCopula(clay_theta_emp, d)
cop_emp_gumb  <- gumbelCopula(gumb_theta_emp, d)
```

A.6 Risk Adjustment Calculation

The portfolio PV reserve distribution is simulated from each copula. Correlated uniform scores are generated from the copula and converted back to monetary PVs using the inverse CDF of each LoB's fitted marginal distribution. The VaR and TVaR RAs are computed at the 75% confidence level.

```
N_COP <- 100000L
set.seed(123)

U_s         <- rCopula(N_COP, cop_obj)
sim_pv      <- u_to_pv(U_s)
total_pv    <- rowSums(sim_pv)

mean_t  <- mean(total_pv)
VaR_75  <- as.numeric(quantile(total_pv, 0.75))
TVaR_75 <- mean(total_pv[total_pv > VaR_75])

RA_VaR  <- VaR_75  - mean_t
RA_TVaR <- TVaR_75 - mean_t
```

A.7 Implied Confidence Level

The implied confidence level is computed by numerically solving for the percentile of the independence copula distribution that corresponds to the copula-based RA.

```
implied_cl <- uniroot(
    f = function(p) {
    as.numeric(quantile(total_indep, p)) -
    mean(total_indep) - RA_target
  },
```

```
  interval = c(0.75, 0.9999),
  tol      = 1e-6
)$root
```

where `total_indep` is the portfolio PV distribution simulated from an independence copula using the same marginal distributions and inverse PIT framework as the six copula specifications.

Appendix B: EIOPA Risk-Free Rates

Table 7.1 shows the EIOPA risk-free rates used for discounting future CF, covering the period from 2026 to 2080. The rates converge to approximately 3.31% from 2065 onwards, reflecting the EIOPA ultimate forward rate assumption.

Table 7.1: EIOPA risk-free rates 2026–2080.

Year	**Rate**	**Year**	**Rate**	**Year**	**Rate**
2026	2.08%	2045	3.21%	2064	3.30%
2027	2.16%	2046	3.22%	2065	3.31%
2028	2.28%	2047	3.23%	2066	3.31%
2029	2.39%	2048	3.24%	2067	3.31%
2030	2.48%	2049	3.25%	2068	3.31%
2031	2.56%	2050	3.26%	2069	3.31%
2032	2.65%	2051	3.27%	2070	3.31%
2033	2.72%	2052	3.27%	2071	3.31%
2034	2.79%	2053	3.28%	2072	3.31%
2035	2.86%	2054	3.28%	2073	3.31%
2036	2.92%	2055	3.28%	2074	3.31%
2037	2.98%	2056	3.29%	2075	3.31%
2038	3.03%	2057	3.29%	2076	3.31%
2039	3.08%	2058	3.29%	2077	3.31%
2040	3.11%	2059	3.30%	2078	3.31%
2041	3.13%	2060	3.30%	2079	3.31%
2042	3.16%	2061	3.30%	2080	3.31%
2043	3.18%	2062	3.30%		
2044	3.19%	2063	3.30%		

Appendix C: Risk Adjustment at Multiple Confidence Levels

Table 7.2: VaR RA at multiple confidence levels

Copula	**Matrix**	**75%**	**90%**	**95%**	**99.5%**
Independence		10,891,661	21,739,795	28,406,608	46,685,109
Gumbel	Empirical	12,162,933	27,071,802	37,374,432	73,775,509
Gumbel	SII	12,454,604	27,376,412	38,281,313	75,683,994
Gaussian	Empirical	13,006,555	25,679,468	33,676,912	56,435,740
Gaussian	SII	13,495,669	26,975,128	35,552,139	59,376,556
Clayton	Empirical	14,023,998	25,797,290	32,988,806	52,111,521
Clayton	SII	14,284,434	26,296,629	33,298,214	52,474,735
Perfect positive dependence		17,108,193	35,699,672	48,216,453	87,661,820

Table 7.3: TVaR Risk Adjustment at multiple confidence levels

Copula	**Matrix**	**75%**	**90%**	**95%**	**99.5%**
Independence		21,702,609	30,642,183	36,532,878	53,771,270
Gumbel	SII	28,701,172	43,329,808	54,460,412	93,094,628
Gumbel	Empirical	28,081,732	42,279,583	53,026,639	90,258,313
Gaussian	SII	27,090,373	38,505,229	46,143,658	68,625,701
Clayton	SII	26,091,661	35,690,164	41,921,968	60,531,122
Gaussian	Empirical	25,830,078	36,605,273	43,829,256	65,748,199
Clayton	Empirical	25,734,472	35,331,177	41,560,452	59,681,973
Perfect positive dependence		36,561,228	53,457,282	65,516,364	105,740,618

Tables 7.2 and 7.3 present the VaR and TVaR RAs at confidence levels of 75%, 90%, 95% and 99.5%. A notable pattern emerges across confidence levels. While the Clayton

copulas produce the highest VaR RA at 75%, the Gumbel copulas overtake them at higher confidence levels, producing substantially larger RAs at 90%, 95% and 99.5%. This is a direct consequence of upper tail dependence. The Gumbel copula concentrates joint adverse outcomes in the right tail, which becomes very apparent at higher confidence levels. At 99.5%, the Empirical Gumbel VaR RA of 73,775,509 is 41.6% higher than the Empirical Clayton RA of 52,111,521. This means that the choice of copula family becomes increasingly important as the confidence level rises.