Stockholms universitet

A Comparison of Different Reinsurance Contracts: Stop-Loss, Excess-of-Loss and Proportional Reinsurance

Anton Pinet Magnusson

Kandidatuppsats i matematisk statistik
Bachelor Thesis in Mathematical Statistics

Kandidatuppsats 2026:4
Matematisk statistik
Maj 2026

www.math.su.se

Matematisk statistik
Matematiska institutionen
Stockholms universitet
106 91 Stockholm

Matematiska institutionen

Mathematical Statistics
Stockholm University
Bachelor Thesis **2026:4**
http://www.math.su.se

A Comparison of Different Reinsurance Contracts: Stop-Loss, Excess-of-Loss and Proportional Reinsurance

Anton Pinet Magnusson*

May 2026

Abstract

This thesis compares proportional reinsurance, stop-loss reinsurance, and per-claim excess-of-loss reinsurance from the cedent's perspective. The comparison is made in a collective risk model where annual losses are simulated using different claim frequency and claim severity distributions. The contracts are calibrated with an equal-risk approach and compared by their expected ceded loss. The results show that stop-loss is generally the most cost-efficient contract at the main 80% retained-risk target. Proportional reinsurance is usually the most expensive alternative, while excess-of-loss becomes more competitive when claim sizes are heavy-tailed.

*Postal address: Mathematical Statistics, Stockholm University, SE-106 91, Sweden. E-mail: anton.p.magnusson123@gmail.com. Supervisor: Leo Levenius and Ola Hössjer.

Abstract

This thesis compares proportional reinsurance, stop-loss reinsurance, and per-claim excess-of-loss reinsurance from the cedent's perspective. The comparison is made in a collective risk model where annual losses are simulated using different claim frequency and claim severity distributions. The contracts are calibrated with an equal-risk approach and compared by their expected ceded loss. The results show that stop-loss is generally the most cost-efficient contract at the main 80% retained-risk target. Proportional reinsurance is usually the most expensive alternative, while excess-of-loss becomes more competitive when claim sizes are heavy-tailed.

Sammanfattning

Detta examensarbete jämför proportionell återförsäkring, stop-loss och per-claim excess-of-loss från cedentens perspektiv. Jämförelsen görs i en kollektiv riskmodell där årliga skadekostnader simuleras med olika modeller för skadefrekvens och skadestorlek. Avtalen kalibreras med en equal-risk-metod och jämförs utifrån den förväntade skadekostnad som cederas till återförsäkraren. Resultaten visar att stop-loss generellt är det mest kostnadseffektiva avtalet vid det huvudsakliga målet om 80% kvarhållen risk. Proportionell återförsäkring är oftast det dyraste alternativet, medan excess-of-loss blir mer konkurrenskraftigt när skadestorlekarna är tungsvansade.

Acknowledgements

ChatGPT has been used to generate code for the simulations, to generate tables and graphs, to search course material and textbooks for relavant information, and to translate some words and formulations to English since this is not my native language.

I want to thank both my supervisors for great engagement, and I would like to express my sincere gratitude to Leo Levenius for his support throughout this thesis. Leo played an important role in helping me find the direction of the project and in shaping the idea that ultimately became this thesis. His thoughtful guidance, insightful feedback, and willingness to discuss ideas have been invaluable. I am especially grateful for having had him as a sounding board during the process, and his encouragement and support have meant a great deal to me.

Notation

- N: number of claims in a year
- X_i: individual claim size for claim i
- $S = \sum_{i=1}^{N} X_i$: annual claim amount total
- R_C: retained annual loss under contract C
- Z_C: ceded annual loss under contract C
- $\mu_C = E[R_C]$: expected retained loss under contract C
- $\sigma_C = \sqrt{\mathrm{Var}(R_C)}$: standard deviation of retained loss
- $\rho_\alpha^{\mathrm{SD}}(R_C) = \mu_C + c_\alpha \sigma_C$: standard-deviation-based monetary risk measure, where $c_\alpha = \Phi^{-1}(\alpha)$ and ϕ is the standard normal distribution
- $\mathrm{VaR}_\alpha(R_C)$: Value at Risk at level α
- $\mathrm{ES}_\alpha(R_C)$: Expected Shortfall at level α
- $E[Z_C]$: expected ceded loss used in this thesis as a cost proxy
- η: target retention ratio
- λ: intensity parameter in the Poisson frequency model
- r, δ: shape and rate parameters in the Poisson–Gamma frequency model
- β: rate parameter in the exponential severity model
- a, b: shape and rate parameters in the gamma severity model
- τ, c: shape and scale parameters in the Weibull severity model
- α, κ: tail index and scale parameter in the Pareto severity model
- K: retention level in stop-loss reinsurance
- D: retention level in per-claim excess-of-loss reinsurance
- p: ceded proportion in proportional reinsurance
- $q = 1 - p$: retained proportion in proportional reinsurance
- $x^+ = \max\{x, 0\}$: positive part used for ceded losses

- M: number of Monte Carlo replications
- Pilot sample and main sample: separate samples used for calibration and final evaluation in the Monte Carlo simulations

Contents

1 Introduction

1.1 Background

Reinsurance is insurance for insurance companies. A direct insurer, often called the *cedent*, transfers part of its risk to a *reinsurer* in exchange for a premium. In this way, the cedent can reduce its exposure to large losses, smooth results over time, release capital, and underwrite larger or more policies than would otherwise be possible (Albrecher et al., 2017, pp. 1–4; Mikosch, 2009, pp. 142–143).

The idea of sharing large risks has a long history. Albrecher et al. (2017) describe a documented marine-related reinsurance contract from 1370 and note that the first specialised reinsurance company was founded in Cologne in 1846 (Albrecher et al., 2017, p. 17). Today, reinsurance is especially important for catastrophic risks, where losses can be large and difficult for a single insurer to absorb (Mikosch, 2009, pp. 142–143).

This thesis focuses on three classical forms of reinsurance: proportional reinsurance, stop-loss, and per-claim excess-of-loss. These contracts differ in how losses are divided between the cedent's retained amount and the reinsurer's ceded amount. Proportional reinsurance cedes a fixed share of the risk, per-claim excess-of-loss covers the part of individual claims above a retention, and stop-loss covers the part of the aggregate loss above a retention (Mikosch, 2009, pp. 142–143; Albrecher et al., 2017, pp. 19–20, 24, 30–31).

1.2 Purpose and research question

The purpose of this thesis is to compare three classical forms of reinsurance from the cedent's perspective: proportional reinsurance, stop-loss, and per-claim excess-of-loss. The comparison is carried out in a stylized collective risk model for annual aggregate loss (Mikosch, 2009, pp. 3–5; Albrecher et al., 2017, p. 189). The research question is: *From the cedent's perspective, which of proportional, stop-loss, and per-claim excess-of-loss reinsurance contracts achieves a given target level of retained annual loss risk at the lowest expected ceded cost, and how does the answer depend on the claim frequency and severity model?* The question consists of two parts. The first concerns the actual ranking of the contracts at a given risk target. The second concerns how this ranking is affected by assumptions regarding claim count and claim sizes.

2 Model framework

2.1 The collective risk model

Let N denote the number of claims during a year and let $X_1, X_2, \ldots, X_N$ be the possible individual claim sizes. The claim sizes are assumed to be independent and identically distributed, non-negative, and independent of N. The annual total claim amount is then defined as

$$S = \sum_{i=1}^{N} X_i,$$

where the empty sum is defined as zero, so that $S = 0$ when $N = 0$.

This is the classical *collective risk model* also called the *compound risk model*. In this model N is the claim frequency component, while the distribution of a generic claim size X is the claim severity component. Thus, the frequency model determines how many claims occur, and the severity model determines how large the individual claims are. When claim arrivals are modelled by a homogeneous Poisson process, the corresponding continuous-time model is the Cramér–Lundberg model, and the aggregate claim amount is a compound sum or compound Poisson sum (Mikosch, 2009, pp. 3–4, 9–12).

Under the assumption of independence, the following relations hold

$$E[S] = E[N]E[X],$$

and

$$\mathrm{Var}(S) = \mathrm{Var}(N)E[X]^2 + E[N]\mathrm{Var}(X).$$

These relations show that both the frequency model and the severity model affect the aggregated annual risk (Albrecher et al., 2017, pp. 189–190).

2.2 Frequency models

2.2.1 Poisson

The first frequency model is the Poisson distribution,

$$N \sim \mathrm{Poisson}(\lambda).$$

Equivalently,

$$\Pr(N = k) = e^{-\lambda}\frac{\lambda^k}{k!}, \qquad k = 0, 1, 2, \ldots.$$

Then

$$E[N] = \lambda, \qquad \mathrm{Var}(N) = \lambda.$$

The Poisson model is the classical starting point for modelling claim counts. It describes a situation in which the variance of the claim count equals the expected value (Mikosch, 2009, pp. 7–8).

2.2.2 Poisson–Gamma

The Poisson model can be too restrictive when claim count data are more variable than the Poisson assumption allows. This situation is called overdispersion and means that

$$\mathrm{Var}(N) > E[N].$$

In insurance applications, overdispersion may arise because the annual claim intensity is affected by unobserved conditions, such as weather, portfolio composition, or other external factors. To allow for this, a mixed Poisson model is used. We assume that there exists a stochastic annual intensity Λ such that

$$\Lambda \sim \mathrm{Gamma}(r, \delta), \qquad N \mid \Lambda \sim \mathrm{Poisson}(\Lambda),$$

where the gamma distribution is parameterized with shape parameter r and rate parameter δ.

By iterated expectation and variance calculation,

$$E[N] = E[\Lambda] = \frac{r}{\delta},$$

and

$$\mathrm{Var}(N) = E[\Lambda] + \mathrm{Var}(\Lambda) = \frac{r}{\delta} + \frac{r}{\delta^2}.$$

Since

$$\mathrm{Var}(\Lambda) = \frac{r}{\delta^2} > 0,$$

the Poisson–Gamma model gives

$$\mathrm{Var}(N) > E[N].$$

The marginal distribution of N is then negative binomial, and the model is suitable as a simple way to represent overdispersion in claim counts (Albrecher et al., 2017, pp. 146–148; Mikosch, 2009, pp. 66–68).

2.3 Severity models

Throughout this section, F_X denotes the distribution function,

$$F_X(x) = \Pr(X \leq x),$$

$\bar{F}_X$ denotes the survival function,

$$\bar{F}_X(x) = \Pr(X > x) = 1 - F_X(x),$$

and f_X denotes the density when a density exists.

2.3.1 Exponential

The first severity model is the exponential distribution,

$$X \sim \mathrm{Exp}(\beta), \qquad \beta > 0.$$

Its distribution function and survival function are

$$F_X(x) = 1 - e^{-\beta x}, \qquad \bar{F}_X(x) = e^{-\beta x}, \qquad x \geq 0.$$

The expected value and variance are

$$E[X] = \frac{1}{\beta}, \qquad \mathrm{Var}(X) = \frac{1}{\beta^2}.$$

The exponential distribution is used here as a simple light-tailed reference model (Mikosch, 2009, p. 96).

2.3.2 Gamma

As a second severity model, the gamma distribution with shape parameter a and rate parameter b is used:

$$X \sim \mathrm{Gamma}(a, b), \qquad f_X(x) = \frac{b^a}{\Gamma(a)} x^{a-1} e^{-bx}, \qquad x > 0.$$

Then

$$E[X] = \frac{a}{b}, \qquad \mathrm{Var}(X) = \frac{a}{b^2}.$$

The gamma distribution is included as a flexible light-tailed alternative to the exponential distribution (Mikosch, 2009, p. 96). With the parameters used below,

$a = 2$ and $b = 2$, it has mean one and variance $1/2$, and is therefore less dispersed than the exponential model with mean one.

2.3.3 Weibull

As the third severity model, the Weibull distribution is used in the form

$$X \sim \text{Weibull}(\tau, c), \qquad \bar{F}_X(x) = \Pr(X > x) = e^{-cx^\tau}, \qquad x \geq 0,$$

where $\tau > 0$ is the shape parameter and $c > 0$ is a positive parameter in this tail parameterization. In the conventional scale-parameter form, the scale is $c^{-1/\tau}$.

The expected value is

$$E[X] = c^{-1/\tau}\Gamma\left(1 + \frac{1}{\tau}\right),$$

and the variance is

$$\text{Var}(X) = c^{-2/\tau}\left[\Gamma\left(1 + \frac{2}{\tau}\right) - \Gamma\left(1 + \frac{1}{\tau}\right)^2\right].$$

For $\tau < 1$, the Weibull survival function decays more slowly than an exponential survival function, so it represents a heavier-tailed case than the exponential model. At the same time, all power moments are finite. The scenario below uses $\tau = 0.7$ to represent this moderately heavy-tailed severity case (Mikosch, 2009, pp. 98, 105).

2.3.4 Pareto

As a heavy-tailed reference model, a Pareto distribution is used:

$$X \sim \text{Pareto}(\alpha, \kappa), \qquad \bar{F}_X(x) = \Pr(X > x) = \left(\frac{\kappa}{\kappa + x}\right)^\alpha, \qquad x \geq 0,$$

where $\alpha > 0$ is the tail index and $\kappa > 0$ is the scale parameter. For $\alpha > 1$,

$$E[X] = \frac{\kappa}{\alpha - 1},$$

and for $\alpha > 2$ the variance is finite and given by

$$\text{Var}(X) = \frac{\alpha\kappa^2}{(\alpha - 1)^2(\alpha - 2)}.$$

Pareto is used here to represent a case in which very large individual claims can play a substantial role in tail risk (Mikosch, 2009, pp. 88, 98; Albrecher et al., 2017, pp. 39–40, 45–46).

2.4 Parameter choice and scenario design

The scenario design is chosen so that the total expected annual claim amount is kept constant across all scenarios. Let

$$\mu_N = E[N], \qquad \mu_X = E[X].$$

Throughout the thesis we set

$$\mu_N = 10, \qquad \mu_X = 1,$$

which gives

$$E[S] = \mu_N \mu_X = 10.$$

In this way, differences between scenarios can be interpreted primarily as differences in dispersion and tail behaviour rather than differences in scale.

For the frequency models, the following parameters are used. In the Poisson case,

$$\lambda = 10,$$

so that

$$E[N] = 10, \qquad \mathrm{Var}(N) = 10.$$

In the Poisson–Gamma case,

$$r = 5, \qquad \delta = \frac{r}{\mu_N} = \frac{5}{10} = 0.5.$$

This gives

$$E[N] = \frac{r}{\delta} = 10,$$

and

$$\mathrm{Var}(N) = \frac{r}{\delta} + \frac{r}{\delta^2} = 10 + 20 = 30.$$

Thus, the Poisson–Gamma model has the same expected claim count as the Poisson model, but a higher variance.

For the severity models, $E[X] = 1$ is maintained in all cases by the following parameter choices:

- Exponential: $\beta = 1$, so that $E[X] = 1$ and $\mathrm{Var}(X) = 1$.
- Gamma: $a = 2$ and $b = a/\mu_X = 2$, so that $E[X] = 1$ and $\mathrm{Var}(X) = 1/2$.
- Weibull: $\tau = 0.7$ and $c = \Gamma\left(1 + \frac{1}{\tau}\right)^{\tau} = \Gamma\left(1 + \frac{1}{0.7}\right)^{0.7} \approx 1.1794$, so that $E[X] = 1$. This gives $\mathrm{Var}(X) \approx 2.1387$.
- Pareto: $\alpha = 2.5$ and $\kappa = \mu_X(\alpha - 1) = 1.5$, so that $E[X] = 1$ and $\mathrm{Var}(X) = 5$.

This yields a total of eight scenarios: two frequency models combined with four severity models. The exponential and gamma cases are light-tailed reference models, the Weibull case with $\tau = 0.7$ represents a moderately heavy-tailed finite-moment model, and the Pareto case represents a more pronounced heavy-tailed situation. This group of scenarios forms the basis for the subsequent Monte Carlo experiments (see section 5).

3 Reinsurance contracts and retained loss

3.1 General decomposition of the claim total

Let C denote a reinsurance contract. Since the model is considered on a one-year horizon, S denotes the aggregate claim total over this period. For each contract, S is divided into two parts

$$S = R_C + Z_C,$$

where R_C is the cedent's retained loss and Z_C is the ceded loss transferred to the reinsurer. This type of decomposition, where the aggregate claim amount is split into a retained part and a reinsured part, is standard in reinsurance modelling (Albrecher et al., 2017, p. 19).

In this thesis, the retained loss R_C is the object of risk measurement. The cost proxy of contract C is defined as the expected ceded loss,

$$\text{Cost}(C) = E[Z_C].$$

This is a pure loss-based cost measure. It does not include safety loadings, expenses, commissions, or market effects, and should therefore not be interpreted as a full reinsurance premium. A full actuarial reinsurance premium typically contains the expected loss together with a safety loading and additional costs (Albrecher et al., 2017, pp. 218–219).

The case with no reinsurance is used as a reference, although it is not a reinsurance contract in the strict sense. It is represented by

$$R_{\text{NR}} = S, \qquad Z_{\text{NR}} = 0,$$

and serves as the baseline in the equal-risk comparisons later in the thesis.

3.2 Proportional reinsurance

Under *proportional* reinsurance, also called quota-share reinsurance, the cedent transfers a fixed proportion p of each claim to the reinsurer. The cedent then retains the proportion $q = 1 - p$, where $0 \leq p \leq 1$. Since the same proportion is applied to the entire claim total, the retained and ceded losses are

$$R_{\text{Prop}} = qS, \qquad Z_{\text{Prop}} = pS.$$

Proportional reinsurance therefore scales the total claim distribution by a constant factor. In particular, R_{Prop} is a scaled version of S, which makes this contract straightforward to analyse and a natural reference in equal-risk comparisons (Albrecher et al., 2017, pp. 19–20; Mikosch, 2009, p. 143).

3.3 Stop-loss reinsurance

Stop-loss (SL) is an aggregate protection. The cedent chooses a retention $K \geq 0$ and bears the claim total up to this level. Any loss above K is borne by the reinsurer. Therefore,

$$R_{\text{SL}} = \min\{S, K\}, \qquad Z_{\text{SL}} = (S - K)^+,$$

where

$$x^+ = \max\{x, 0\}.$$

Stop-loss acts directly on the aggregate claim total and thereby provides protection against unfavourable claim outcomes as a whole. The version studied here is an unlimited stop-loss contract, meaning that there is no upper cap on the reinsurer's liability. In practice, stop-loss contracts may also include an upper limit, but such capped variants are not considered in this thesis (Albrecher et al., 2017, pp. 30–31; Mikosch, 2009, p. 143).

3.4 Per-claim excess-of-loss

Under per-claim *excess-of-loss*, an individual retention $D \geq 0$ is chosen. The cedent bears each individual claim up to D, while the reinsurer covers the part of the claim that exceeds D. For claim i, the retained part is therefore $\min\{X_i, D\}$ and the ceded part is $(X_i - D)^+$. Aggregating over all claims gives

$$R_{\text{XL}} = \sum_{i=1}^{N} \min\{X_i, D\}, \qquad Z_{\text{XL}} = \sum_{i=1}^{N} (X_i - D)^+.$$

In contrast to stop-loss, XL acts on each individual claim and not directly on the aggregate claim total. It therefore provides protection against very large individual claims, but does not in itself protect against the occurrence of many claims below the retention. In practice, XL contracts are often written with a finite layer, but this thesis considers the unlimited per-claim form above (Albrecher et al., 2017, pp. 24, 30; Mikosch, 2009, p. 143).

3.5 Why these three contracts?

The three selected types of contracts represent three different mechanisms of risk transfer: fixed proportional sharing, an aggregate cap, and a claim-by-claim cap. They are therefore well suited to the research question, which compares how different contract forms perform when the retained risk is held at the same target level. Other contract forms, such as surplus reinsurance, layered XL contracts, and reinstatement clauses, are excluded in order to keep the analysis focused and comparable (Albrecher et al., 2017, pp. 21, 24, 27).

4 Risk measures and pricing

4.1 Starting point

In order to compare reinsurance contracts, both a measure of the cedent's retained risk and a clear cost concept are required. In this thesis, risk is measured through the retained loss R_C, while cost is measured through the ceded loss Z_C. Since the research question concerns which contract achieves a given risk target at the lowest cost, the risk measures and the cost concept must be defined before the numerical analysis.

4.2 Expected value and standard deviation

The most basic measure of retained loss is the expected value

$$\mu_C = \mathbb{E}[R_C].$$

This describes the average loss borne by the cedent under contract C. To measure dispersion around this expected value, the standard deviation $\sigma_C = \sqrt{\mathrm{Var}(R_C)}$ is used.

4.3 A standard deviation-based monetary risk measure

As a simple monetary risk measure, this thesis also uses

$$\rho_\alpha^{\mathrm{SD}}(R_C) = \mu_C + c_\alpha \sigma_C, \qquad c_\alpha = \Phi^{-1}(\alpha),$$

where Φ is the cumulative distribution function of the standard normal distribution. The superscript SD is only used as a label to indicate that the loading is based on the standard deviation.

This measure is closely related to the standard deviation premium principle, where the expected loss is increased by a loading proportional to the standard deviation (Mikosch, 2009, pp. 79–80; Albrecher et al., 2017, p. 220). If R_C were normally distributed, then

$$\mu_C + \Phi^{-1}(\alpha)\sigma_C$$

would equal the upper α-quantile of R_C. If R_C is only approximately normal, the same expression can be interpreted as a normal approximation to that quantile. In this thesis, $\rho_\alpha^{\mathrm{SD}}$ is used as a transparent intermediate between a pure dispersion measure, such as σ_C, and tail-based measures, such as Value at Risk and Expected

Shortfall.

4.4 Value at Risk

For a confidence level $\alpha \in (0, 1)$, *value at risk* is defined as the left α-quantile

$$\mathrm{VaR}_\alpha(R_C) = \inf\{x \in \mathbb{R} : F_{R_C}(x) \geq \alpha\},$$

where F_{R_C} is the cumulative distribution function of R_C. Thus, $\mathrm{VaR}_\alpha(R_C)$ specifies a level that the retained loss exceeds with probability at most $1 - \alpha$. VaR is therefore a natural measure when the focus lies on a high but specific quantile level (Lindskog, 2026, pp. 62–64; Albrecher et al., 2017, p. 128).

4.5 Expected Shortfall

As a tail measure, *expected shortfall* is also used:

$$\mathrm{ES}_\alpha(R_C) = \frac{1}{1-\alpha} \int_\alpha^1 \mathrm{VaR}_u(R_C)\, du.$$

For continuous distributions, this can be written as the conditional expected value of the retained loss given that the loss exceeds its VaRα. In contrast to VaR, ES therefore takes the tail above the quantile level into account and not only a single point of the distribution (Lindskog, 2026, pp. 64–65; Albrecher et al., 2017, pp. 128, 226–227).

4.6 Choice of confidence level

In this thesis, we set $\alpha = 0.995$ for all measures described above. This choice reflects that the analysis concerns serious outcomes, but without going so far into the tail that numerical estimation becomes unnecessarily costly. It is also a standard benchmark in European solvency regulation: under Solvency II, the required capital is based on Value at Risk at level 0.995 over a one-year horizon (Albrecher et al., 2017, p. 226; Lindskog, 2026, p. 65). The use of 0.995 should not be interpreted as a full Solvency II calculation, but it gives an actuarially recognizable severe-risk level. A lower level such as 0.95 would describe less extreme outcomes and would be less aligned with this solvency-motivated perspective.

4.7 Positive homogeneity

A risk measure or premium principle ρ for the random variable X is *positively homogeneous* if, for every constant $q > 0$,

$$\rho(qX) = q\rho(X).$$

This means that scaling a loss by a positive constant scales the measured risk by the same constant. The retained-risk measures used in the main equal-risk analysis—the standard deviation, the standard deviation principle, the Value-at-Risk, and the Expected Shortfall—are all positively homogeneous. For example, if

$$\rho_\alpha^{\mathrm{SD}}(X) = \mathbb{E}[X] + c_\alpha \,\mathrm{SD}(X),$$

then

$$\rho_\alpha^{\mathrm{SD}}(qX) = q\mathbb{E}[X] + c_\alpha q \,\mathrm{SD}(X) = q\rho_\alpha^{\mathrm{SD}}(X).$$

By contrast, the variance principle is not positively homogeneous. For

$$\rho_\alpha^{\mathrm{Var}}(X) = \mathbb{E}[X] + \alpha \,\mathrm{Var}(X),$$

we have

$$\rho_\alpha^{\mathrm{Var}}(qX) = q\mathbb{E}[X] + \alpha q^2 \,\mathrm{Var}(X) \neq q\rho_\alpha^{\mathrm{Var}}(X)$$

(Lindskog, 2026, pp. 62–65; Albrecher et al., 2017, pp. 219–220).

4.8 The concept of cost and premium for the cedent

The quantity used as cost in this thesis is the expected ceded loss, $\mathbb{E}[Z_C]$. This quantity is simple to interpret and directly comparable between contracts. It can also be interpreted as the fair premium, that is, the expected compensation paid by the reinsurer under the contract.

In practice, reinsurance premiums are set higher than the pure expected payout by means of some form of loading. A simple classical premium principle is the expected value principle

$$\Pi_C^{(\theta)} = (1+\theta)\mathbb{E}[Z_C], \qquad \theta \geq 0.$$

If the same loading factor θ were used for all contracts, the ranking between them in terms of cost would be preserved and a comparison based on $\mathbb{E}[Z_C]$ or $\Pi_C^{(\theta)}$ would be equivalent. In this thesis, $\mathbb{E}[Z_C]$ is therefore used as a cost proxy. Any contract specific loadings are left outside the analysis.

4.9 Comparison rule

The main comparison in this thesis is an equal-risk comparison. For a chosen retained-risk measure ρ and a target retention ratio $\eta \in (0, 1]$, the contract parameter is selected so that

$$\rho(R_C) = \eta\rho(S),$$

up to negligible rounding error. Here η is the intended fraction of the unreinsured risk that should remain with the cedent. For example, $\eta = 0.70$ means that the retained risk is targeted at 70% of the no-reinsurance risk.

An alternative comparison rule is an equal-cost comparison, where the expected ceded loss is fixed and contracts are instead compared by the retained risk they produce. This approach is not used in the main analysis of this thesis, but examples of this is provided in Appendix.

The contract that achieves the target retained risk at the lowest expected ceded loss is seen as the most cost-effective for the cedent under the given risk measure.

5 Monte Carlo methodology

This chapter describes how the numerical experiments are performed. *Monte Carlo* simulations are used to estimate empirical distributions of the retained loss R_C and the ceded loss Z_C under the given contracts. This method is chosen because there are no simple closed-form analytical expressions for these distributions in the collective risk model when frequency and severity models are combined with non-proportional reinsurance contracts. Monte Carlo methods are commonly used in such situations, where the total claim amount distribution can be simulated even when exact calculation is difficult (Mikosch, 2009, pp. 130–132; Albrecher et al., 2017, pp. 273–277).

The main analysis is an equal-risk comparison. For every frequency-severity scenario of the total claims S, risk measure ρ, and target retention ratio η, the contract parameters are first chosen on a training sample so that the retained risk is as close as possible to the target risk $\eta\rho(S)$. The same parameters are then evaluated on an independent test sample, and the contracts C are compared by their expected ceded loss $\mathbb{E}[Z_C]$.

5.1 Simulation setup

5.1.1 What is simulated

The analysis is carried out on a one-year horizon. For each scenario, the simulation first generates the annual claim count N, then the individual claim sizes $X_1, \dots, X_N$, and finally the gross annual claim amount

$$S = \sum_{i=1}^{N} X_i.$$

After this, the contract parameter is calibrated according to the equal-risk rule, and the retained and ceded losses are computed for each contract.

For a given frequency-severity scenario of the total claims S, and a given contract C, the central quantities are the retained loss R_C and the ceded loss Z_C. In each simulated year the identity

$$S = R_C + Z_C$$

holds. From the simulated samples, the following quantities are estimated:

$$\mu_C, \qquad \sigma_C, \qquad \rho_\alpha^{\mathrm{SD}}(R_C), \qquad \mathrm{VaR}_\alpha(R_C), \qquad \mathrm{ES}_\alpha(R_C), \qquad \mathbb{E}[Z_C].$$

5.1.2 Training and test samples

To separate calibration from evaluation, two independent simulation samples are used in each frequency-severity scenario. The training sample is used to select the contract parameters. The test sample is used for the results reported in section 6.

This separation is important because the research question is an equal-risk question, meaning that the contracts are to be compared at the same level of retained risk. If the parameters are both chosen and evaluated on the same simulated data, the comparison can become overly optimistic, especially for tail-based risk measures. By separating the training and test samples, a cleaner assessment is obtained of how well the calibrated contract actually achieves the target risk.

5.1.3 Simulation sample sizes

Let M denote the number of simulated years in a sample. In the final simulation design, the same sample sizes are used for all scenarios:

$$M_{\text{train}} = 400\,000, \qquad M_{\text{test}} = 400\,000.$$

Thus, each frequency–severity scenario is calibrated on 400 000 simulated years and evaluated on an independent test sample of 400 000 simulated years.

The same sample sizes are used across all scenarios to make the comparisons as uniform as possible. Large samples are used because the analysis includes $\text{VaR}_{0.995}$ and $\text{ES}_{0.995}$, where only a small fraction of simulated observations contributes directly to the tail estimation. This is especially relevant for heavy-tailed cases and rare-event probabilities, where crude Monte Carlo estimates may require many replications (Mikosch, 2009, pp. 131–132; Albrecher et al., 2017, pp. 276–277).

5.1.4 Target retention ratios and primary target

In the equal-risk comparison, the following target retention ratios are used:

$$\eta \in \{0.95, 0.90, 0.85, 0.80, 0.75, 0.70\}.$$

For a chosen risk measure ρ, the target risk on the training sample is defined as

$$r_0^{\text{train}} = \eta\, \hat{\rho}^{\text{train}}(S),$$

where $\hat{\rho}^{\text{train}}(S)$ is the estimated gross no-reinsurance risk on the training sample.

The main comparison in the results section is performed at

$$\eta = 0.80.$$

The other target ratios are included to show how the cost comparison changes when the retained-risk target is made milder or more restrictive.

5.1.5 Reproducibility and common random numbers

All simulations are governed by a fixed master seed. In the final implementation, one global seed,

$$\texttt{MASTER\_SEED} = 19981118,$$

is set once at the beginning of the simulation. The training and test samples are then generated sequentially from this random number stream. This makes the experiment reproducible without using separate scenario-specific seed formulas.

In addition, within each given frequency-severity scenario the same simulated claim years are used for all contracts. This means that proportional reinsurance, stop-loss, and per-claim excess-of-loss are evaluated on the same underlying claim experience year by year. This corresponds to the use of common random numbers and reduces Monte Carlo noise in the contract comparisons without changing the marginal distribution of any contract outcome.

5.2 Simulation of annual claims

5.2.1 Frequency models

Let $m = 1, \dots, M$ index the simulated years. Under the Poisson frequency model, the claim count in simulated year m is generated in R by

$$N^{(m)} \sim \text{Poisson}(\lambda),$$

using `rpois`. Under the Poisson–Gamma frequency model, the mixed Poisson representation is used. For each simulated year m, a stochastic intensity is first generated using `rgamma`,

$$\Lambda^{(m)} \sim \Gamma(r, \delta),$$

where r is the shape parameter and δ is the rate parameter. Conditional on this intensity, the claim count is then generated using `rpois`,

$$N^{(m)} \mid \Lambda^{(m)} \sim \text{Poisson}(\Lambda^{(m)}).$$

5.2.2 Severity models

Given $N^{(m)}$, the individual claim sizes

$$X_1^{(m)}, \dots, X_{N^{(m)}}^{(m)}$$

are generated from the severity model corresponding to the scenario. In the implementation, R's random variate generators `rexp`, `rgamma`, and `rweibull` are used for the Exponential, Gamma, and Weibull distributions, respectively. For the Pareto distribution, inverse transform sampling is used. If $U \sim U(0,1)$, then

$$X = \kappa \left((1-U)^{-1/\alpha} - 1 \right)$$

has the Pareto distribution used in this thesis. This follows by solving the Pareto distribution function for its inverse quantile function (Lindskog, 2026, pp. 4, 10).

In the Pareto case, U is generated using `runif` and clipped away from the exact endpoints to avoid numerical boundary problems. Once the claim sizes have been simulated, the gross loss in year m is calculated as

$$S^{(m)} = \sum_{i=1}^{N^{(m)}} X_i^{(m)}.$$

5.2.3 Calculation of retained and ceded loss

Each simulated claim year is then used to calculate the retained and ceded loss under the three types of contracts included in the research question. The actual calibration of the parameters p, K, and D, that were introduced in Section 3.2-3.4, is described in Section 5.4 below.

With no reinsurance, the baseline is

$$R_{\text{NR}}^{(m)} = S^{(m)}, \qquad Z_{\text{NR}}^{(m)} = 0.$$

Under proportional reinsurance with ceded proportion p and retained proportion $q = 1 - p$,

$$R_{\text{Prop}}^{(m)} = qS^{(m)}, \qquad Z_{\text{Prop}}^{(m)} = pS^{(m)}.$$

Under stop-loss with retention K,

$$R_{\text{SL}}^{(m)} = \min\{S^{(m)}, K\}, \qquad Z_{\text{SL}}^{(m)} = (S^{(m)} - K)^+.$$

Under per-claim excess-of-loss with retention D,

$$R_{\mathrm{XL}}^{(m)} = \sum_{i=1}^{N^{(m)}} \min\{X_i^{(m)}, D\}, \qquad Z_{\mathrm{XL}}^{(m)} = \sum_{i=1}^{N^{(m)}} (X_i^{(m)} - D)^+.$$

5.3 Estimation of risk measures and cost proxy

After M simulated years, where M is either M_{train} in the calibration step or M_{test} in the evaluation step, we obtain for each type of contract C, for training and test data, the samples

$$\{R_C^{(m)}\}_{m=1}^M, \qquad \{Z_C^{(m)}\}_{m=1}^M.$$

The expected retained loss and the expected ceded loss are estimated by the sample means

$$\hat{\mu}_C = \bar{R}_C = \frac{1}{M}\sum_{m=1}^M R_C^{(m)}, \qquad \widehat{\mathbb{E}[Z_C]} = \bar{Z}_C = \frac{1}{M}\sum_{m=1}^M Z_C^{(m)}.$$

The retained standard deviation is estimated by

$$\hat{\sigma}_C = \sqrt{\frac{1}{M-1}\sum_{m=1}^M \left(R_C^{(m)} - \bar{R}_C\right)^2}.$$

The standard deviation-based monetary risk measure is then estimated as

$$\widehat{\rho}_\alpha^{\mathrm{SD}}(R_C) = \bar{R}_C + c_\alpha \hat{\sigma}_C, \qquad c_\alpha = \Phi^{-1}(\alpha).$$

Let $R_{C,(1)} \le \cdots \le R_{C,(M)}$ denote the ordered simulated retained losses and let $k = \lceil \alpha M \rceil$. The value at risk is estimated as

$$\widehat{\mathrm{VaR}}_\alpha(R_C) = R_{C,(k)}.$$

The empirical distribution function has empirical quantile function

$$\widehat{F}_M^{\leftarrow}(u) = R_{C,(j)} \quad \text{for} \quad \frac{j-1}{M} < u \le \frac{j}{M}.$$

Inserting this step function into the integral definition

$$\widehat{\mathrm{ES}}_\alpha(R_C) = \frac{1}{1-\alpha}\int_\alpha^1 \widehat{F}_M^{\leftarrow}(u)\,du$$

gives an estimate

$$\widehat{\mathrm{ES}}_\alpha(R_C) = \frac{1}{1-\alpha}\left[\left(\frac{k}{M}-\alpha\right)R_{C,(k)} + \frac{1}{M}\sum_{j=k+1}^{M} R_{C,(j)}\right]$$

of the expected shortfall. The first term appears because α may fall inside the interval on which the empirical quantile function equals $R_{C,(k)}$. If αM is an integer, this term is zero. This estimator is the empirical version of the integrated-quantile definition of Expected Shortfall and is the estimator used in the simulation (Lindskog, 2026, pp. 25–26; Albrecher et al., 2017, pp. 226–227).

5.4 Calibration of parameters and comparison rule

5.4.1 Equal-risk as the main method

The main experiment addresses the research question by holding retained risk fixed and then comparing cost this is the method used in the thesis to answer the research question. For each scenario, retained-risk measure ρ, and target risk fraction η, the contract parameters are calibrated on the training sample. The training-sample target risk level is

$$T^{\text{train}}_{\rho,\eta} = \eta\,\widehat{\rho}^{\text{train}}(S),$$

where $\widehat{\rho}^{\text{train}}(S)$ is the estimated no-reinsurance risk on the training sample.

For proportional reinsurance, the calibration is exact for the retained-risk measures used in the equal-risk analysis. In the implementation, these are standard deviation, the standard deviation principle, Value-at-Risk, and Expected Shortfall. Since

$$R_{\text{Prop}} = qS$$

and these retained-risk measures are positively homogeneous, as discussed in Section 4.7,

$$\rho(R_{\text{Prop}}) = \rho(qS) = q\rho(S).$$

Thus, to attain the target

$$\rho(R_{\text{Prop}}) = \eta\rho(S),$$

we set

$$q = \eta, \qquad p = 1-\eta.$$

For stop-loss reinsurance, two different calibration methods are used. For VaR_α, the exact relation

$$\mathrm{VaR}_\alpha(\min\{S, K\}) = \min\{\mathrm{VaR}_\alpha(S), K\}$$

is used. This holds because the stop-loss retained loss $\min\{S, K\}$ is the gross loss capped at K. If the gross α-quantile is below K, the cap does not affect the quantile, while if the gross α-quantile is above K, the retained loss cannot exceed K. Therefore, if the target is

$$\mathrm{VaR}_\alpha(R_{\mathrm{SL}}) = \eta\, \mathrm{VaR}_\alpha(S),$$

we choose

$$K = \eta\, \widehat{\mathrm{VaR}}_\alpha^{\mathrm{train}}(S).$$

For σ_C, $\rho_\alpha^{\mathrm{SD}}(R_C)$, and $\mathrm{ES}_\alpha(R_C)$, the stop-loss retention K is found numerically on the training sample. For a stop-loss contract with retention K, the retained annual loss in simulated year m is

$$R_{\mathrm{SL}}^{(m)}(K) = \min\{S^{(m)}, K\}.$$

The calibration problem is therefore to find the largest value of K such that

$$\widehat{\rho}^{\mathrm{train}}(R_{\mathrm{SL}}(K)) \le T_\eta = T_{\rho,\eta}^{\mathrm{train}}.$$

The function

$$K \mapsto \widehat{\rho}^{\mathrm{train}}(\min\{S, K\})$$

is non-decreasing for the retained-risk measures considered here. Hence the retention can be found through the bisection method. The algorithm starts with a lower bound

$$K_L = 0$$

and an upper bound

$$K_U = \max_m S^{(m)}.$$

At each step, the midpoint

$$K_M = \frac{K_L + K_U}{2}$$

is evaluated. If

$$\widehat{\rho}^{\mathrm{train}}(R_{\mathrm{SL}}(K_M)) \le T_\eta,$$

then the target is still attained and the retention can be increased. In that case, K_L is replaced by K_M. Otherwise, the retained risk is above the target and K_U is replaced by K_M. After the iterations, K_L is selected. This gives the largest retention found by the search that still attains the target. This choice is used because the expected ceded loss

$$\widehat{\mathbb{E}}^{\text{train}}[(S-K)_+]$$

decreases as K increases.

For per-claim excess-of-loss, the retained annual loss for deductible D is

$$R_{\text{XL}}^{(m)}(D) = \sum_{i=1}^{N^{(m)}} \min\{X_i^{(m)}, D\},$$

and the ceded annual loss is

$$Z_{\text{XL}}^{(m)}(D) = \sum_{i=1}^{N^{(m)}} (X_i^{(m)} - D)_+.$$

In this simulation design, no closed-form inversion of

$$D \mapsto \widehat{\rho}^{\text{train}}(R_{\text{XL}}(D))$$

is used. Instead, D is chosen by an empirical grid search.

The grid is constructed from the individual $M = M_{\text{train}}$ claim sizes $R_{\{XL\}}^{(m)}(D)$ in the training sample. Let

$$\widehat{F}_X^{\leftarrow}(p)$$

denote the empirical p-quantile of these M claim sizes. Candidate deductibles D are obtained from equally spaced quantile levels in the intervals

$$[0, 0.90], \quad [0.90, 0.99], \quad [0.99, 0.999], \quad [0.999, 0.9999], \quad [0.9999, 0.99999],$$

together with $D = 0$, the maximum observed claim size, and a value just above the maximum observed claim size. The grid is therefore denser in the upper tail, where changes in the XL deductible are most relevant for tail-risk measures. This construction is based on empirical quantiles, as described by Lindskog (2026, pp. 25–26).

For each risk measure ρ and candidate D, the retained losses $R_{\text{XL}}^{(m)}(D)$, the ceded losses $Z_{\text{XL}}^{(m)}(D)$, the retained risk, and the expected ceded loss are calculated on the training sample. The candidate deductibles are then ordered in increasing D. Since a larger deductible means that the cedent retains more of each claim, the retained-risk curve is treated as non-decreasing in the calibration. To avoid selecting a deductible because of small empirical irregularities in the simulated risk curve, the estimated curve is replaced by its cumulative maximum. More precisely, if

$$0 \leq D_1 < \cdots < D_L$$

are the ordered grid values and

$$\widehat{r}_\ell = \widehat{\rho}^{\text{train}}(R_{\text{XL}}(D_\ell)),$$

then the adjusted risk curve is

$$\widetilde{r}_\ell = \max_{j \leq \ell} \widehat{r}_j.$$

The selected deductible is the largest D_ℓ for which

$$\widetilde{r}_\ell \leq T_\eta.$$

This gives the largest deductible on the grid that still attains the target. Since the expected ceded loss decreases as D increases, this also selects the cheapest deducible attaining XL candidate on the grid. If no grid point attains the target, the nearest grid point is used, and the resulting row is later assessed using the test-sample target-gap rule.

Thus, the stop-loss calibration is a root-search problem on a monotone empirical risk curve, whereas the XL calibration is a discrete grid search over empirically chosen deductible values. The stop-loss and excess-of-loss contract forms used here are standard reinsurance forms; see Lindskog (2026, pp. 39–40) and Albrecher et al. (2017, pp. 24, 30). The use of simulation-based sample averages for the retained risk and expected ceded loss follows the Monte Carlo principle described by Albrecher et al. (2017, pp. 273–274).

5.4.2 Evaluation on the test sample

Once the parameters have been selected on the training sample, they are applied unchanged to the independent test sample. All tables and figures in the results section are based on this test sample.

The purpose of the test sample is to check whether the calibrated contract still reaches the intended retained-risk level on independent simulated data. For this reason, the realised retained risk on the test sample is compared with the test-sample equivalent target risk

$$T^{\text{test}}_{\rho,\eta} = \eta\, \widehat{\rho}^{\text{test}}(S),$$

where $\widehat{\rho}^{\text{test}}(S)$ is the estimated no-reinsurance risk on the test sample. The relative test gap is defined as

$$\text{gap}^{\text{test}} = \frac{\widehat{\rho}^{\text{test}}(R_C) - T^{\text{test}}_{\rho,\eta}}{T^{\text{test}}_{\rho,\eta}}.$$

A contract row is considered usable for ranking only if

$$|\text{gap}^{\text{test}}| \leq 0.05.$$

This 5% rule is not a statistical test. It is a practical filter used to avoid ranking contracts in cases where the equal-risk condition is not achieved closely enough.

5.5 Numerical checks and reproducibility

To ensure that differences between types of contracts are not caused by coding errors or numerical instability, several types of checks are performed.

First, the identity

$$S^{(m)} = R_C^{(m)} + Z_C^{(m)}$$

is checked for every simulated year and every contract. Second, the exact formulas for proportional reinsurance, stop-loss, and XL are verified against the simulated retained and ceded losses.

The training and test target gaps are also stored. This makes it possible to see whether a parameter fits the training sample but misses the target on the independent test sample. In addition, the sample is divided into $B = 20$ blocks, and VaR_α and ES_α are calculated separately in each block. The standard deviation of these block estimates is used as a simple diagnostic for Monte Carlo variation in the tail estimates.

Finally, all metadata required for reproduction are saved: frequency-severity scenario, master seed, sample size, risk measure, target retention ratio, contract parameters, and estimated results. This allows the reported tables and figures to be regenerated from the saved output files.

5.6 Summary of the Experiment

The final Monte Carlo setup can be summarised in four steps.

1. Select a frequency–severity scenario and simulate an independent training sample and test sample, each with 400 000 simulated years.
2. Calibrate proportional reinsurance, stop-loss, and per-claim excess-of-loss on the training sample so that the chosen retained-risk measure reaches the target retention ratio η of the gross no-reinsurance risk.

3. Evaluate the same parameters on the test sample and compute $\hat{\sigma}_C$, $\hat{\rho}^{\mathrm{SD}}_{0.995}(R_C)$, $\widehat{\mathrm{VaR}}_{0.995}(R_C)$, $\widehat{\mathrm{ES}}_{0.995}(R_C)$, and $\widehat{\mathbb{E}[Z_C]}$.

4. Rank the contracts only in those cases where the relative test-sample target gap is sufficiently small.

This setup makes the comparison directly tailored to the research question.

6 Results

This section presents the numerical results of the equal-risk comparison between proportional reinsurance, stop-loss reinsurance, and per-claim excess-of-loss reinsurance. The purpose is to examine how the ranking of the contracts is affected by the frequency model, the severity model, the chosen retained-risk measure, and the target risk fraction. Throughout the section, the theoretical cost proxy for contract C is the expected ceded loss $c_C = \mathbb{E}[Z_C]$ as defined in Section 4. In the numerical results, this quantity is estimated on the independent test sample by

$$\hat{c}_C^{\text{test}} = \widehat{\mathbb{E}}^{\text{test}}[Z_C] = \frac{1}{M_{\text{test}}} \sum_{m=1}^{M_{\text{test}}} Z_C^{(m)}.$$

Thus, reported costs refer to the Monte Carlo estimate $\widehat{c}_C^{\text{test}}$.

6.1 Experimental design

The equal-risk comparison is conducted across eight frequency-severity scenarios obtained by combining two frequency models with four severity models. The frequency models are Poisson, and Poisson–gamma. The severity models are exponential, gamma, Weibull, and Pareto. The scenario design keeps the theoretical expected gross loss fixed at

$$\mathbb{E}[S] = \mathbb{E}[N]\mathbb{E}[X] = 10$$

in all scenarios. Small deviations from 10 in the reported simulated means,

$$\widehat{\mathbb{E}}^{\text{test}}[S] = \frac{1}{M_{\text{test}}} \sum_{m=1}^{M_{\text{test}}} S^{(m)},$$

are therefore Monte Carlo sampling variation rather than differences in the theoretical scale.

The main comparison is performed at the target risk fraction $\eta = 0.80$. This means that, for a given retained-risk measure ρ, the contract parameter is calibrated on the training sample so that

$$\widehat{\rho}^{\text{train}}(R_C) = \eta\, \widehat{\rho}^{\text{train}}(S).$$

The calibrated parameter is then applied unchanged to the independent test sample, where the reported retained risk and cost are evaluated.

6.2 Numerical validation

The numerical identity checks are mainly coding checks. They verify that the simulated retained and ceded losses satisfy the defining contract identities. In the final run, the maximum absolute error in

$$S^{(m)} = R_C^{(m)} + Z_C^{(m)}$$

was zero at the reported precision, and the corresponding formula checks for proportional reinsurance, stop-loss, and per-claim excess-of-loss also had zero maximum absolute error. The numerical differences discussed below are therefore not caused by implementation errors.

More important for interpretation is whether the equal-risk target is achieved on the independent test sample. Across the full equal-risk grid, the target check is applied to all non-baseline reinsurance rows. There are 8 frequency-severity scenarios, 4 retained-risk measures, 6 target retention ratios, and 3 reinsurance contracts, namely proportional reinsurance, stop-loss, and per-claim excess-of-loss. This gives

$$8 \cdot 4 \cdot 6 \cdot 3 = 576$$

non-baseline contract rows. The no-reinsurance case is excluded from this count because it is only a baseline and no contract parameter is calibrated for it. Of these 576 rows, 557 satisfy the 5% rule. The 19 rows that fall outside the rule are all XL rows and all have negative test gaps, meaning that the retained risk is slightly below the target rather than above it.

At the primary target $\eta = 0.80$, there are 32 combinations of frequency-severity scenarios and retained-risk measures. In 28 of these, all three contracts are usable for ranking according to the 5% rule. In the remaining four combinations, the non-usable row is always XL. The affected cases are Poisson +exponential under σ_C and $\rho^{\mathrm{SD}}_{0.995}$, Poisson–gamma +exponential under σ_C, and Poisson–gamma +exponential under $\mathrm{ES}_{0.995}$. The ranking statements below are therefore based on the usable rows, while these four XL cases are interpreted with caution.

6.3 Gross risk without reinsurance

Before comparing the contracts, it is useful to examine the gross risk profile without reinsurance. Table 1 reports the simulated no-reinsurance risk measures for the eight frequency-severity scenarios.

Under the Poisson frequency model, gamma produces the lowest dispersion, followed by exponential. Weibull yields higher risk measures, while Pareto gives by far the largest tail measures. For example, under Poisson frequency, $\mathrm{VaR}_{0.995}(S)$ increases from 21.766 in the gamma case and 24.172 in the exponential case to 29.815 in the Weibull case and 41.852 in the Pareto case. The corresponding $\mathrm{ES}_{0.995}(S)$ reaches 62.736 in the Poisson + Pareto scenario.

When the Poisson frequency model is replaced by Poisson–gamma, the reported risk measures increase in all severity cases. In the exponential case, for instance, $\sigma(S)$ rises from 4.470 to 6.338, $\mathrm{VaR}_{0.995}(S)$ from 24.172 to 32.296, and $\mathrm{ES}_{0.995}(S)$ from 26.610 to 36.386. Thus, overdispersion in the claim count amplifies unfavourable outcomes even though the theoretical expected gross loss is held fixed.

Table 1: Gross risk without reinsurance for the eight scenarios.

Scenario	$\mathbb{E}[S]$	$\sigma(S)$	$\rho^{\mathrm{SD}}_{0.995}(S)$	$\mathrm{VaR}_{0.995}(S)$	$\mathrm{ES}_{0.995}(S)$
Poisson +exponential	9.995	4.470	21.511	24.172	26.610
Poisson + gamma($a = 2$)	10.000	3.876	19.982	21.766	23.605
Poisson + Weibull($\tau = 0.7$)	9.995	5.600	24.420	29.815	33.615
Poisson + Pareto($\alpha = 2.5$)	9.990	7.495	29.294	41.852	62.736
Poisson–gamma +exponential	10.010	6.338	26.334	32.296	36.386
Poisson–gamma + gamma($a = 2$)	10.015	5.931	25.294	30.516	34.364
Poisson–gamma + Weibull($\tau = 0.7$)	9.990	7.160	28.432	36.203	41.421
Poisson–gamma + Pareto($\alpha = 2.5$)	9.990	8.922	32.974	46.342	67.000

6.4 Main equal-risk comparison at the 80% target

The central question is which contract achieves a given level of retained risk at the lowest expected ceded loss. At the target retention ratio of 80%, proportional reinsurance provides a natural benchmark. Since

$$R_{\mathrm{Prop}} = qS,$$

The proportional calibration is therefore exact with

$$q = 0.80, \qquad p = 0.20,$$

which yields

$$\mathbb{E}[Z_{\mathrm{Prop}}] \approx 2$$

in all scenarios and for all risk measures.

Table 2 reports the expected ceded loss for all three contracts at the primary target $\eta = 0.8$. The final column indicates whether the XL row is usable under the 5% target-gap rule. Proportional reinsurance and stop-loss are usable in all rows shown.

The main result is that stop-loss is the cheapest usable contract in every scenario-risk combination at the primary target. In the 28 combinations where all three contracts are usable, stop-loss is always cheaper than both XL and proportional reinsurance. In the four combinations where XL is not usable, the observed cost ordering is still stop-loss, XL, and proportional reinsurance, but these XL rows are not used for a full ranking.

Table 2: Expected ceded loss $\mathbb{E}[Z_C]$ at the equal-risk target $\eta = 0.80$.

Scenario	Risk measure	Prop	SL	XL	XL usable
Poisson +exponential	σ_C	1.999	0.482	1.447	FALSE
Poisson +exponential	$\rho^{\text{SD}}_{0.995}$	1.999	0.845	1.897	FALSE
Poisson +exponential	$\text{VaR}_{0.995}$	1.999	0.086	1.447	TRUE
Poisson +exponential	$\text{ES}_{0.995}$	1.999	0.043	1.447	TRUE
Poisson + gamma($a = 2$)	σ_C	2.000	0.441	1.486	TRUE
Poisson + gamma($a = 2$)	$\rho^{\text{SD}}_{0.995}$	2.000	0.833	1.795	TRUE
Poisson + gamma($a = 2$)	$\text{VaR}_{0.995}$	2.000	0.084	1.486	TRUE
Poisson + gamma($a = 2$)	$\text{ES}_{0.995}$	2.000	0.041	1.486	TRUE
Poisson + Weibull($\tau = 0.7$)	σ_C	1.999	0.515	0.940	TRUE
Poisson + Weibull($\tau = 0.7$)	$\rho^{\text{SD}}_{0.995}$	1.999	0.826	1.263	TRUE
Poisson + Weibull($\tau = 0.7$)	$\text{VaR}_{0.995}$	1.999	0.090	0.775	TRUE
Poisson + Weibull($\tau = 0.7$)	$\text{ES}_{0.995}$	1.999	0.042	0.661	TRUE
Poisson + Pareto($\alpha = 2.5$)	σ_C	1.998	0.173	0.197	TRUE
Poisson + Pareto($\alpha = 2.5$)	$\rho^{\text{SD}}_{0.995}$	1.998	0.377	0.430	TRUE
Poisson + Pareto($\alpha = 2.5$)	$\text{VaR}_{0.995}$	1.998	0.167	0.197	TRUE
Poisson + Pareto($\alpha = 2.5$)	$\text{ES}_{0.995}$	1.998	0.061	0.067	TRUE
Poisson–gamma +exponential	σ_C	2.002	0.593	1.904	FALSE
Poisson–gamma +exponential	$\rho^{\text{SD}}_{0.995}$	2.002	0.888	1.904	TRUE
Poisson–gamma +exponential	$\text{VaR}_{0.995}$	2.002	0.100	1.904	TRUE
Poisson–gamma +exponential	$\text{ES}_{0.995}$	2.002	0.047	1.904	FALSE
Poisson–gamma + gamma($a = 2$)	σ_C	2.003	0.573	1.800	TRUE
Poisson–gamma + gamma($a = 2$)	$\rho^{\text{SD}}_{0.995}$	2.003	0.883	2.117	TRUE
Poisson–gamma + gamma($a = 2$)	$\text{VaR}_{0.995}$	2.003	0.095	1.800	TRUE
Poisson–gamma + gamma($a = 2$)	$\text{ES}_{0.995}$	2.003	0.046	1.800	TRUE
Poisson–gamma + Weibull($\tau = 0.7$)	σ_C	1.998	0.612	1.260	TRUE
Poisson–gamma + Weibull($\tau = 0.7$)	$\rho^{\text{SD}}_{0.995}$	1.998	0.880	1.515	TRUE
Poisson–gamma + Weibull($\tau = 0.7$)	$\text{VaR}_{0.995}$	1.998	0.103	1.208	TRUE
Poisson–gamma + Weibull($\tau = 0.7$)	$\text{ES}_{0.995}$	1.998	0.048	1.156	TRUE
Poisson–gamma + Pareto($\alpha = 2.5$)	σ_C	1.998	0.350	0.446	TRUE
Poisson–gamma + Pareto($\alpha = 2.5$)	$\rho^{\text{SD}}_{0.995}$	1.998	0.580	0.840	TRUE
Poisson–gamma + Pareto($\alpha = 2.5$)	$\text{VaR}_{0.995}$	1.998	0.175	0.283	TRUE
Poisson–gamma + Pareto($\alpha = 2.5$)	$\text{ES}_{0.995}$	1.998	0.069	0.076	TRUE

For the exponential and gamma scenarios, stop-loss is clearly the cheapest contract. For example, in Poisson +exponential, the expected ceded loss under σ_C is 0.482 for stop-loss, compared with 1.447 for XL and 1.999 for proportional reinsurance. Under $\text{VaR}_{0.995}$, the corresponding costs are 0.086, 1.447, and 1.999.

The same qualitative ordering appears for the Weibull scenarios. In Poisson +

Weibull, the expected ceded loss is 0.515 for stop-loss, 0.940 for XL, and 1.999 for proportional reinsurance under σ_C. Under $\mathrm{VaR}_{0.995}$, the corresponding costs are 0.090, 0.775, and 1.999.

For the Pareto scenarios, stop-loss remains the cheapest at the primary target, but XL becomes much more competitive. In Poisson + Pareto($\alpha = 2.5$), the cost under σ_C is 0.173 for stop-loss and 0.197 for XL. Under $\mathrm{ES}_{0.995}$, the difference is even smaller: 0.061 for stop-loss and 0.067 for XL. In Poisson–gamma + Pareto($\alpha = 2.5$), the corresponding $\mathrm{ES}_{0.995}$ costs are 0.069 for stop-loss and 0.076 for XL. Thus, the heavy-tailed Pareto case is where XL comes closest to stop-loss, even though it does not overtake it at $\eta = 0.80$.

Figures 1 and 2 show the same pattern under Poisson frequency for σ_C and $\mathrm{VaR}_{0.995}$. The figures should be read together with Table 2, since the table also records whether the XL row is usable for ranking.

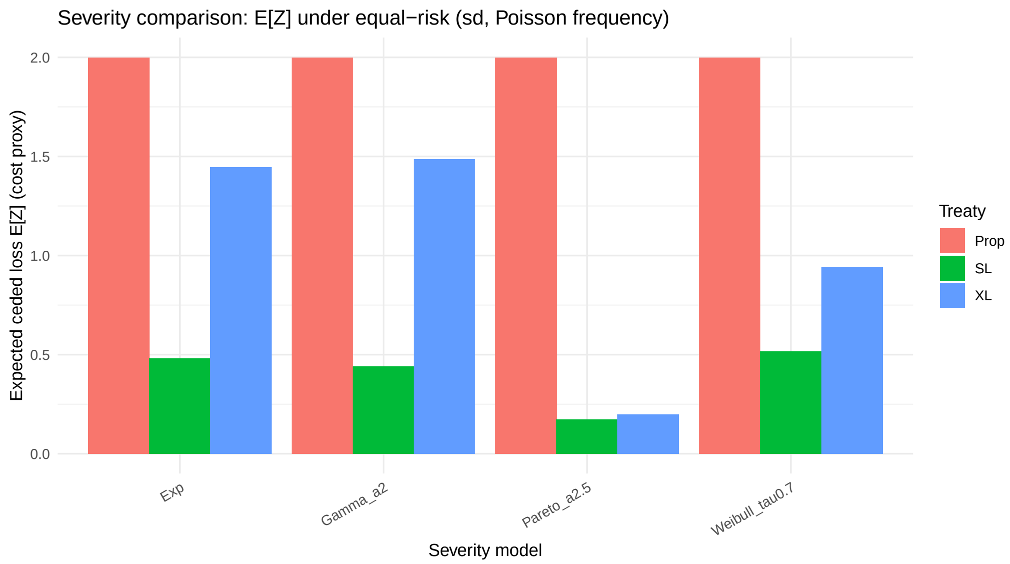

Figure 1: Equal-risk cost comparison across severity models under Poisson frequency, with σ_C as the target risk measure.

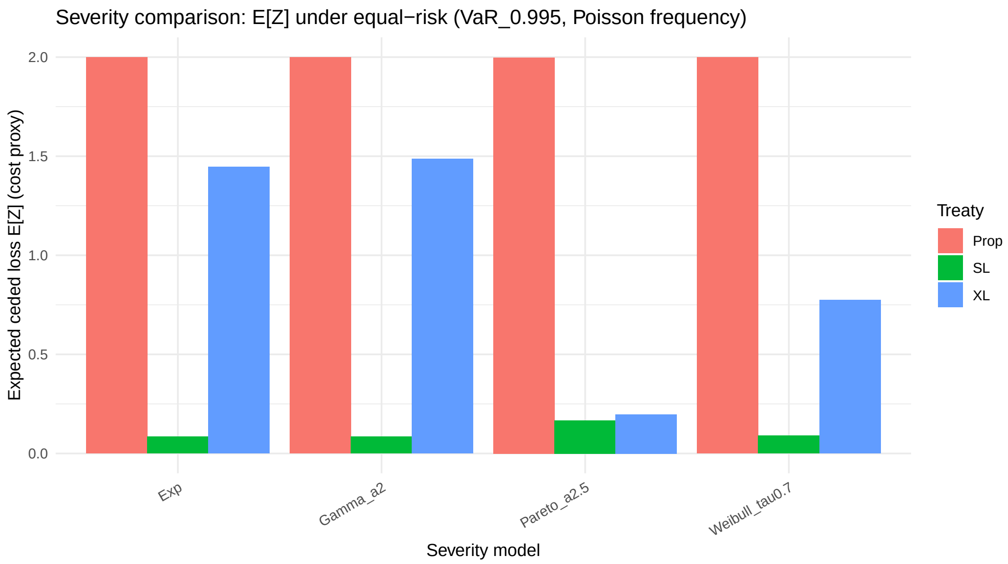

Figure 2: Equal-risk cost comparison across severity models under Poisson frequency, with $\mathrm{VaR}_{0.995}(R_C)$ as the target risk measure.

6.5 The effect of the target retention ratio

The primary comparison above uses $\eta = 0.80$, but the simulation also includes the target ratios

$$\eta \in \{0.95, 0.90, 0.85, 0.80, 0.75, 0.70\}.$$

This makes it possible to examine how the cost ordering changes when the required risk reduction becomes milder or more restrictive. Figures 3 and 4 show the cost curves for the baseline Poisson +exponential scenario. They illustrate three general features. First, proportional reinsurance increases approximately linearly in cost as η decreases, because $p = 1 - \eta$. Second, the stop-loss cost remains below the costs of both proportional reinsurance and XL over the displayed target range. Third, XL lies between stop-loss and proportional reinsurance in this light-tailed baseline case. Thus, varying η does not change the ranking in the baseline scenario, but it makes the cost difference between the contracts more visible.

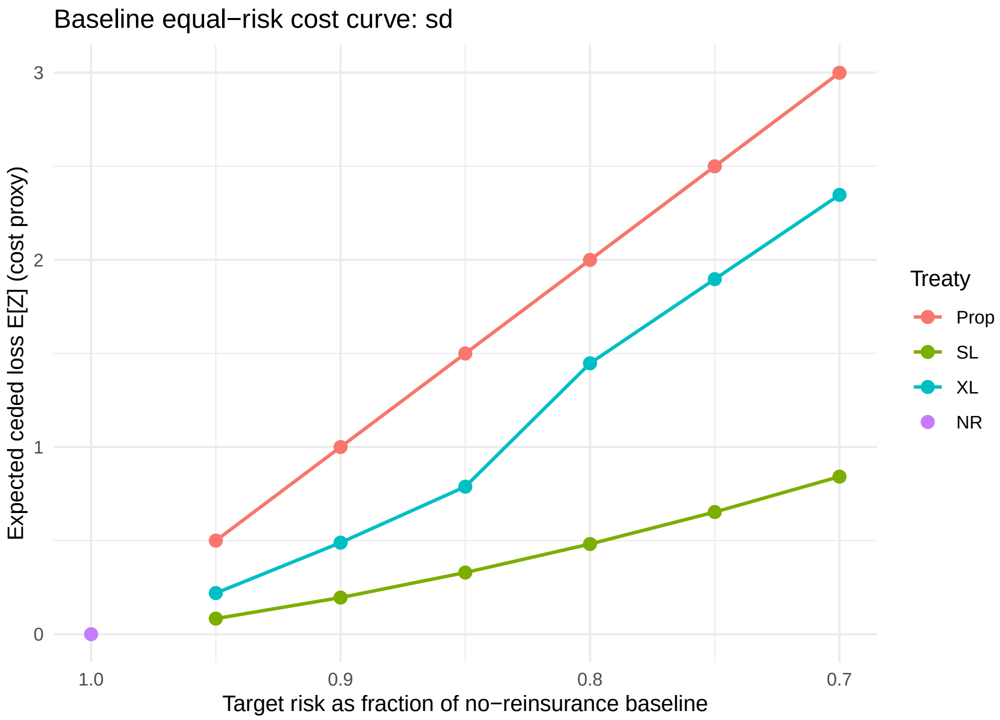

Figure 3: Equal-risk cost curves for the Poisson +exponential scenario, with σ_C as the target risk measure.

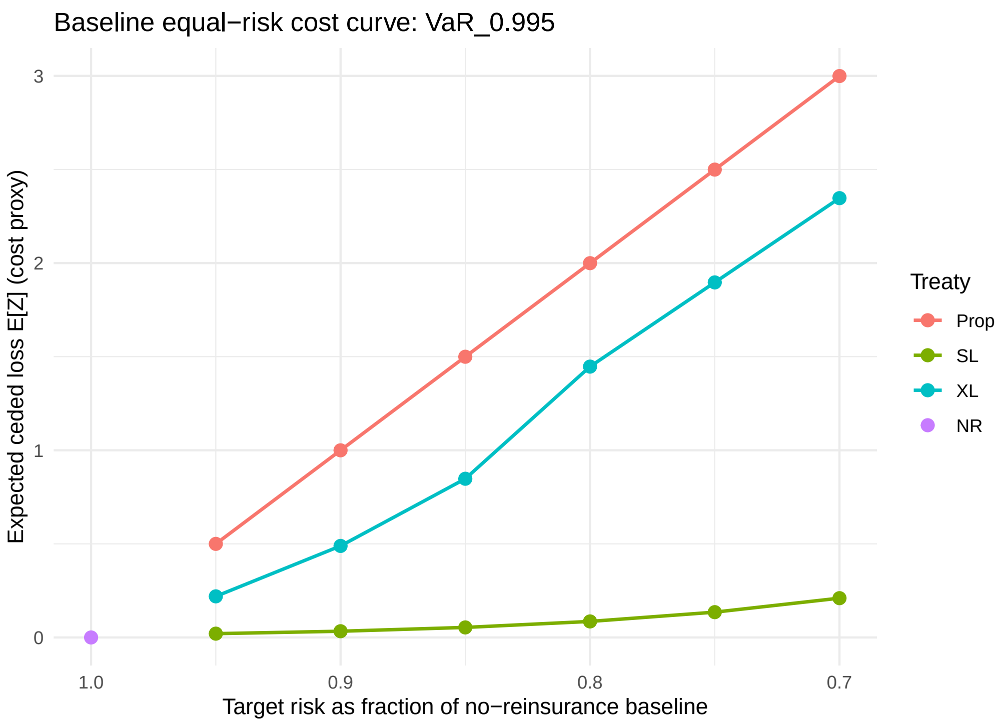

Figure 4: Equal-risk cost curves in the Poisson +exponential scenario, with $\mathrm{VaR}_{0.995}(R_C)$ as the target risk measure.

The most important changes across η occur in the Pareto scenarios. Table 3 reports all cases in which XL is the cheapest usable contract in the equal-risk results. These switches only occur in the Pareto scenarios and mainly under $\mathrm{VaR}_{0.995}$. For the Poisson + Pareto scenario, XL is slightly cheaper than stop-loss under $\mathrm{VaR}_{0.995}$ already at $\eta = 0.85$, and clearly cheaper at $\eta = 0.90$ and $\eta = 0.95$. In the Poisson–gamma + Pareto scenario, XL becomes cheaper than stop-loss under $\mathrm{VaR}_{0.995}$ at $\eta = 0.90$ and $\eta = 0.95$. There is also one very small switch under $\mathrm{ES}_{0.995}$ at $\eta = 0.90$ in the Poisson–gamma + Pareto scenario, where the costs are almost identical.

Table 3: Cases where XL is the cheapest usable contract when the target retention ratio varies.

Scenario	Risk measure	η	SL	XL
Poisson + Pareto($\alpha = 2.5$)	$\mathrm{VaR}_{0.995}$	0.85	0.147	0.146
Poisson + Pareto($\alpha = 2.5$)	$\mathrm{VaR}_{0.995}$	0.90	0.130	0.116
Poisson + Pareto($\alpha = 2.5$)	$\mathrm{VaR}_{0.995}$	0.95	0.116	0.094
Poisson–gamma + Pareto($\alpha = 2.5$)	$\mathrm{VaR}_{0.995}$	0.90	0.132	0.119
Poisson–gamma + Pareto($\alpha = 2.5$)	$\mathrm{VaR}_{0.995}$	0.95	0.116	0.090
Poisson–gamma + Pareto($\alpha = 2.5$)	$\mathrm{ES}_{0.995}$	0.90	0.037	0.036

Figures 5 and 6 show the corresponding $\mathrm{VaR}_{0.995}$-based cost curves for the two Pareto scenarios. These figures explain why the ranking at $\eta = 0.80$ should not be interpreted as universal across all target levels. At stronger risk reductions, such as $\eta = 0.70$ and $\eta = 0.75$, stop-loss is still cheaper. At milder risk reductions, however, XL becomes cheaper under $\mathrm{VaR}_{0.995}$, because it can remove the most extreme individual claims with a relatively small expected ceded loss.

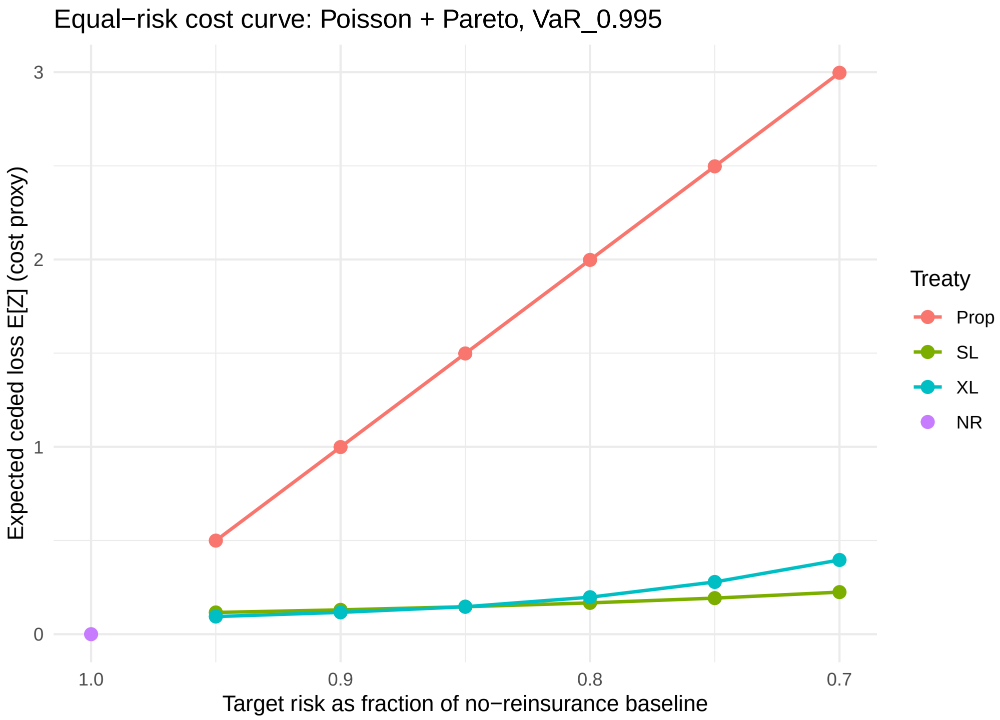

Figure 5: Equal-risk cost curves in the Poisson + Pareto scenario, with $\mathrm{VaR}_{0.995}(R_C)$ as the target risk measure.

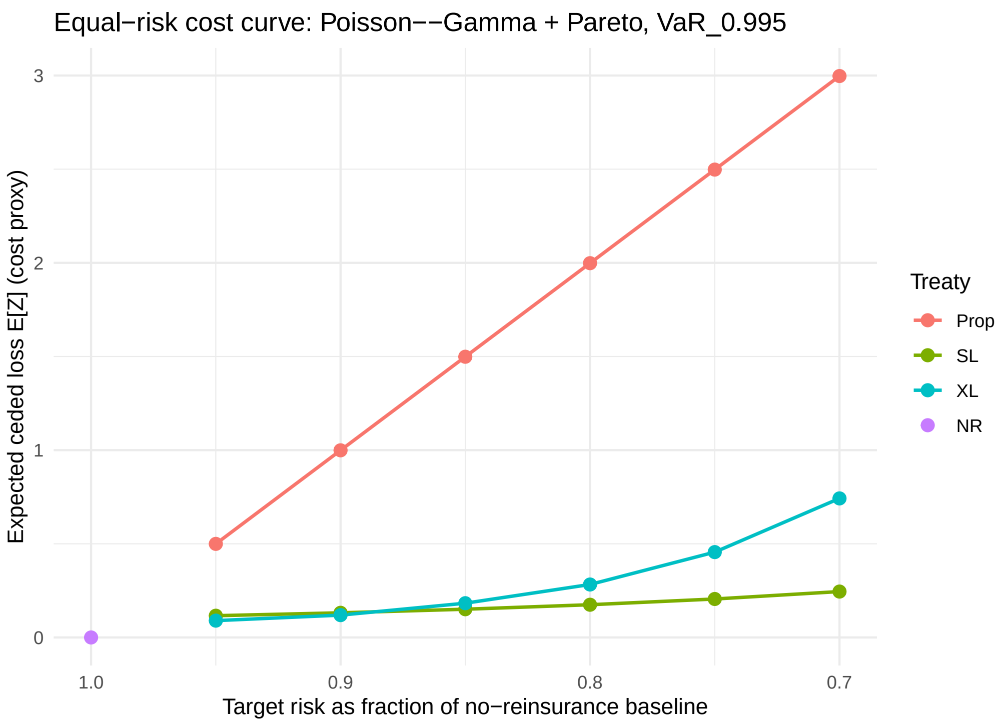

Figure 6: Equal-risk cost curves in the Poisson–gamma + Pareto scenario, with $\mathrm{VaR}_{0.995}(R_C)$ as the target risk measure.

6.6 The effect of the severity distribution

The clearest pattern in the results is that the tail behaviour of the severity distribution has a strong effect on how close the cost of XL comes to that of stop-loss. As shown in Table 2 and Figures 1–2, stop-loss is far below both XL and proportional reinsurance for the exponential and gamma cases. In these lighter-tailed scenarios, XL is often less expensive than proportional reinsurance, but it is still clearly more expensive than stop-loss.

The Weibull scenario with $\tau = 0.7$ is intermediate. XL becomes less expensive than in the exponential and gamma cases, especially for tail-based measures, but stop-loss still remains the cheapest contract at the primary target. For example, under Poisson frequency and $\mathrm{ES}_{0.995}$, XL costs 0.661 compared with 0.042 for stop-loss and 1.999 for proportional reinsurance.

The Pareto scenario gives the strongest change in relative performance. At the primary target $\eta = 0.80$, stop-loss remains slightly cheaper under all four retained-risk measures. However, Table 3 and Figures 5–6 show that XL can become the cheapest usable contract when the target is milder and the retained risk is measured by $\mathrm{VaR}_{0.995}$. The tail of the severity distribution therefore matters not only for the level of the cost but also for whether the preferred contract can change as η varies.

6.7 The effect of frequency variation

Switching from the Poisson to the Poisson–gamma frequency model increases the gross risk in all severity scenarios, as shown in Table 1. The frequency model also changes the absolute cost levels in Table 2, but it changes the ranking less than the severity distribution does.

For the exponential case under σ_C, the stop-loss cost rises from 0.482 under Poisson frequency to 0.593 under Poisson–gamma frequency. XL rises from 1.447 to 1.904. In the Weibull case under σ_C, stop-loss rises from 0.515 to 0.612, while XL rises from 0.940 to 1.260. Thus, frequency overdispersion generally increases the cost required to reach the same retained-risk target.

In the Pareto scenarios, frequency overdispersion also increases the gross tail risk, but it does not change the main ranking at the primary target. Under $\mathrm{VaR}_{0.995}$, for example, stop-loss costs 0.167 and XL 0.197 in the Poisson + Pareto scenario, while the corresponding costs in the Poisson–gamma + Pareto scenario are 0.175 and 0.283. At $\eta = 0.80$, the frequency model therefore affects the magnitude of the costs but not the conclusion that stop-loss is cheapest contract. At milder targets, especially $\eta = 0.90$ and $\eta = 0.95$, XL becomes the cheapest contract under $\mathrm{VaR}_{0.995}$ in both

Pareto frequency settings.

6.8 Summary of empirical results

Overall, the results show that stop-loss is the cheapest contract at the primary target retention ratio $\eta = 0.80$ in every scenario-risk combination where a full ranking is permitted by the 5% maximum gap rule. For the 28 combinations where all three reinsurance contracts are usable, stop-loss is always the cheapest. In the four remaining combinations, XL is not usable for full ranking, but the observed cost ordering is still stop-loss, followed by XL, and finally proportional reinsurance.

The severity distribution has the strongest effect on the relative performance of the contracts. In the exponential and gamma scenarios, stop-loss is clearly cheaper than XL and proportional reinsurance. In the Weibull scenario, XL becomes more competitive but still remains above stop-loss at the primary target. In the Pareto scenario, XL comes closest to stop-loss, especially under $\mathrm{ES}_{0.995}$, but stop-loss remains marginally cheaper at $\eta = 0.80$.

The additional target-ratio analysis shows that the primary-target conclusion should be read together with the value of η. At the primary target $\eta = 0.80$, stop-loss is the most cost-effective contract in the studied scenario set. When the target is milder, especially $\eta = 0.90$ and $\eta = 0.95$, XL becomes the cheapest usable contract under $\mathrm{VaR}_{0.995}$ in the Pareto scenarios. This is the clearest evidence in the simulation that heavy-tailed individual claims can make per-claim excess-of-loss more attractive than stop-loss for some risk targets.

The frequency model mainly affects the absolute cost levels. Moving from Poisson to Poisson–gamma increases gross risk and usually increases the expected ceded loss required to reach the same retained-risk target. However, within the studied scenarios, the change from Poisson to Poisson–gamma has less impact on the contract ranking than the change from light-tailed to heavy-tailed Pareto severity.

The most robust empirical conclusion is therefore that stop-loss is the dominant contract at the primary equal-risk target, while XL becomes increasingly attractive when very large individual claims dominate the tail risk and the retained-risk target is less restrictive, corresponding to a larger value of η.

7 Discussion

This chapter interprets the results in relation to the research question and discusses what drives the different outcomes for the studied reinsurance contracts. The discussion is conducted within the equal-risk and model framework chosen earlier. This means that the cheapest contract here refers to the contract that yields the lowest expected ceded loss, $\mathbb{E}[Z_C]$, since this has been the cost proxy used throughout the numerical analysis. It should therefore not be interpreted as a market premium, as previously discussed.

7.1 Discussion in relation to the research question

Our research question was the following:

> *From the cedent's perspective, which of proportional, stop-loss, and per-claim excess-of-loss reinsurance contracts achieves a given target level of retained annual loss risk at the lowest expected ceded cost, and how does the answer depend on the claim frequency and severity model?*

Our results showed that the question can be answered clearly at the primary target retention ratio $\eta = 0.80$, however, to stay true to the results we have to formulate this answer with precision. When the equal-risk comparison is used, and we only compare the candidates that fulfill the 5%-rule maximum gap-rule, stop-loss is the cheapest for all 28 combinations of frequency-severity scenario and risk measure in which all three reinsurance contracts could be fully ranked. In the four remaining exponential cases the XL contract failed the 5% target-gap tolerance, but the observed cost ordering is nevertheless the same. Stop-loss is cheapest, followed by XL, which in turn is followed by proportional reinsurance. The most fair answer to the research question with the given primary target retention ratio would be to confirm that the results point to that the most cost-effective contract among the ranked scenario-risk combinations would be stop-loss whereas for the unranked ones this still seems to be the case.

At the same time the results show that the answer to the research question is not invariant to the model choice. It is in particular the severity model's tail behavior that affects the contracts' value for the cedent. With the primary target retention ration $\eta = 0.8$, stop-loss remains the cheapest even for the Pareto scenarios, however XL becomes clearly more competitive than in the exponential, gamma, and weibull scenarios. When the retained risk is measured with $\mathrm{VaR}_{0.995}(R_C)$ and the target risk fraction is set to 0.90 or 0.95, then XL actually becomes the cheapest contract in the Pareto scenario under both frequency models. There is also a smaller switch at $\eta = 0.85$ in the Poisson–Pareto case and one very small switch under $ES_{0.995}(R_C)$ at $\eta = 0.90$ in the Poisson–Gamma–Pareto case. The answer to the research question is therefore not that one contract is universally best in all environments but rather that stop-loss dominates for the studied combination of frequency-severity scenario and risk measure at the primary target retention ratio, while XL becomes the most cost-effective contract in certain heavy-tailed, quantile-based comparisons.

7.2 Why stop-loss wins in most cases

The most striking result is the position of stop-loss in the light-tailed and moderately heavy-tailed scenarios. This is a reasonable result intuitively. Since the research question concerns retained annual loss risk, in other words risk measured at the aggregate annual level, stop-loss acts directly on annual aggregate claims by capping retained loss at level K. The contract therefore addresses precisely the quantity the comparison concerns. Consequently, when the risk measure is defined on the annual aggregate retained loss, an aggregate protection is naturally well suited to the problem.

This is very apparent in the exponential, gamma, and weibull scenarios. In these cases the retained risk is reduced most cost-effectively by protecting against large annual aggregate losses, rather than by proportionally sharing each individual claim or by only capping individual claim sizes. Although proportional reinsurance does provide exact calibration, it achieves this by scaling down the entire loss distribution, including the small and medium-sized claims that do not primarily drive the higher retained risk measures. It is therefore natural that proportional reinsurance becomes one of the more expensive solutions when cost is expressed as $c_C = \mathbb{E}[Z_C]$, and in many comparisons it is the most expensive contract.

Nevertheless, the fact that XL is often cheaper than proportional reinsurance in these scenarios demonstrates that the individual truncation of large claims has a clear value. However, the results also show that this is generally not sufficient to outperform stop-loss when the retained risk is primarily driven by the sum of many claims during the year. In precisely these environments, stop-loss emerges as the contract form that makes the most efficient use of ceded cost.

7.3 Why XL becomes more promising with heavy tails

The Pareto scenarios show that the choice between contracts becomes more subtle when very large individual claims become more prominent. In these cases, the gross risk measures rise sharply, in particular $\mathrm{VaR}_{0.995}(S)$ and $\mathrm{ES}_{0.995}(S)$, and the retained risk is influenced to a greater extent by a few extreme observations rather than by the sum of many moderate claims. In these cases per-claim excess-of-loss naturally becomes more relevant.

This explains why XL closes in on stop-loss in the Pareto scenarios and in some cases even surpasses it. At the primary target retention ratio $\eta = 0.80$, stop-loss remains the cheapest contract, yet the difference relative to XL becomes very small under $\mathrm{ES}_{0.995}(R_C)$. In a practical sense the two contracts are almost equivalent within the studied model. When the retained risk is instead measured by $\mathrm{VaR}_{0.995}(R_C)$ and the target reduction is made milder, XL becomes the cheapest usable contract for both Pareto scenarios at $\eta = 0.90$ and $\eta = 0.95$. In addition, XL becomes the cheapest usable contract at $\eta = 0.85$ in the Poisson + Pareto scenario, and one very small switch appears under $\mathrm{ES}_{0.995}(R_C)$ at $\eta = 0.90$ in the Poisson–Gamma + Pareto scenario. This pattern is consistent with the fact that a quantile-based risk measure in a heavy-tailed environment is largely driven by a limited number of very large individual claims. Capping precisely these claims can then be more cost-effective than imposing an aggregate cap on the entire annual loss.

At the same time this result should not be overinterpreted. It is not the heavy-tailed severity in itself that automatically makes XL the best contract. At the primary target retention ratio $\eta = 0.80$, stop-loss remains the cheapest option even in the Pareto scenarios. The results therefore do not suggest that XL should generally be preferred in heavy-tailed environments. Rather they show that XL becomes relatively more attractive only when two conditions coincide, that being when the tail is sufficiently heavy and the retained risk is measured by a quantile-based risk measure under a less aggressive target reduction(i.e. a larger η).

7.4 The role of the risk measure

A central lesson from the study is that the choice of reinsurance contract cannot be separated from the choice of risk measure. The results show that the ranking of the contracts is most stable under σ_C and $\rho^{\mathrm{SD}}_{0.995}(R_C)$, where stop-loss clearly dominates across almost the entire scenario group. Under $\mathrm{ES}_{0.995}(R_C)$ the same main pattern persists, although the difference between stop-loss and XL becomes very small in the Pareto scenarios. The most pronounced shifts in contract ranking occur under $\mathrm{VaR}_{0.995}(R_C)$, where XL becomes the cheapest usable contract for both Pareto scenarios at $\eta = 0.90$ and $\eta = 0.95$. In addition, XL becomes the cheapest usable contract at $\eta = 0.85$ in the Poisson + Pareto scenario, and one very small switch appears under $\mathrm{ES}_{0.995}(R_C)$ at $\eta = 0.90$ in the Poisson–Gamma + Pareto scenario.

This illustrates that the four retained-risk measures capture different aspects of risk. The standard deviation and the standard deviation-based monetary risk measure respond to the overall dispersion in the retained loss. VaR, on the other hand, focuses on a single high quantile, whereas ES accounts for the entire extreme tail above the quantile level. And so the fact that different contracts become most attractive under different risk measures is not a methodological issue, rather it constitutes an important part of the answer to the research question. Which contract is the "best" depends on which aspect of retained risk is being prioritised.

This also implies that the results should be interpreted with caution. The cedent whose primary objective is to stabilise the variation in the annual aggregate loss receives stronger support for stop-loss than the cedent who is mainly concerned with very large individual losses in the upper quantile. An overarching contribution of the analysis is therefore precisely to make this dependence structure between risk measure and contract choice explicit.

7.5 The role of frequency overdispersion

The transition from the Poisson to the Poisson–Gamma frequency model increases the aggregate gross risk as expected, across all severity scenarios. The higher variance in claim count leads to more extreme unfavourable years even though the theoretical expected annual gross loss $E(S)$ is held fixed. This is evident for both the standard deviations and the higher quantiles and tail measures.

By contrast, frequency overdispersion affects the ranking of the contracts considerably less than the severity model does. Within the studied scenario group, overdispersion primarily changes the absolute cost levels rather than which contract is the most cost-effective. This holds particularly at the primary target retention ratio $\eta = 0.80$, where stop-loss remains the cheapest in all fully rankable cases under both frequency models.

This conclusion should however be understood as valid only within the framework of the studied grid, meaning the chosen parameters for the frequency overdispersion. Only one level of overdispersion was used in the Poisson–Gamma model. The results therefore show that the severity tail matters more than this particular change in frequency variation, but they do not establish any general rule for all conceivable frequency models or all levels of overdispersion.

7.6 Methodological assessment and limitations

Methodologically, the study has several strengths. The most important is that it is directly tailored to the research question. The equal-risk design forces a comparison in which the contracts are evaluated at the same level of retained risk rather than at the same contract parameters or some other more ambiguous measure. Furthermore, by first selecting the parameters on a training sample and then evaluating them on an independent test sample, target attainment can be assessed without confounding calibration and results. The use of common random numbers within each scenario also improves the comparability between the contracts by reducing Monte Carlo noise in the differences between their costs.

Another strength is that the study does not treat all numerical outcomes as equally reliable. By explicitly defining a 5% maximum gap rule for usable ranking, it becomes clear which rows can be used for stronger conclusions and which should be interpreted with greater caution. This enhances the internal rigour of the analysis and reduces the risk that low costs are misinterpreted as contractual advantages when the target risk has in fact not been achieved. Across the full equal-risk grid, 557 out of 576 non-baseline contract rows satisfy this rule. The remaining rows are all XL rows with negative test gaps, meaning that the retained risk of the cedent is slightly below the target value rather than above it.

At the same time, there are clear limitations. First, cost is defined as $\mathbb{E}[Z_C]$, i.e. the expected ceded loss. This constitutes a transparent and comparable cost proxy, but it is not equivalent to an actual reinsurance premium in practice. A market premium is also influenced by loadings, expenses, commissions, contract terms, reinstatement clauses, and market conditions. The results should therefore be interpreted as answering the question of the lowest expected ceded loss at a given retained-risk target, rather than as a complete answer to which treaty would certainly be the cheapest under real-world market pricing (Albrecher et al., 2017, pp. 218–220, 228–229).

Secondly, the model environment is deliberately stylized in order to directly address the research question. The frequency and severity models are assumed independent, the analysis is conducted on a one-year horizon, and only three contract forms are studied. Other aspects such as parameter uncertainty, claim inflation, dependence between claims, layered structures, reinstatements, and alternatives to annual aggregate risk as the decision criterion are not included in the main experiment. The results should therefore be understood as a well-defined answer within the collective risk model, rather than as a universal reinsurance theorem.

7.7 Practical implications from the cedent's perspective

From the cedent's perspective, the results provide relatively clear practical guidance. If the risk target concerns the annual aggregate retained loss distribution and the portfolio is not dominated by extremely heavy individual claims, the results strongly support stop-loss as the most cost-effective contract form. This holds particularly when the focus lies on the standard deviation or a standard deviation-based monetary risk measure.

By contrast, if the portfolio is characterized by very large individual claims and the relevant risk assessment is quantile-based, XL becomes a clearly more relevant alternative. The results do not indicate that proportional reinsurance is a cost-effective first-choice option in the studied environment when the objective is expressed as annual aggregate retained loss risk. It may however still be practically justified for other reasons not modelled here, such as simplicity, commission structures, capital management, strategic relationships with reinsurers, pricing and loading considerations, or situations where broad protection is desired also for smaller and medium-sized claims (Albrecher et al., 2017, pp. 2–4, 218–220, 228–229).

A key practical takeaway is therefore that the choice of reinsurance contract should not be made without an explicit risk measure. The same portfolio of insurance costumers can lead to different recommendations depending on whether the decision-maker focuses on overall annual dispersion, a high quantile, or average tail loss. The study thereby demonstrates not only which treaty is often the most cost-effective, but also why it is essential to first define precisely which aspect of risk one actually wishes to reduce.

7.8 Concluding discussion

The results show that the research question does not admit a fully contract-independent answer. However, within the studied scenario group the main pattern is clear: stop-loss is the most cost-effective contract at the primary target $\eta = 0.8$ and remains the cheapest even for the Pareto scenarios. At the same time, XL becomes progressively more attractive as the severity tail grows heavier and, in certain usable Pareto cases, emerges as the cheapest contract when the comparison is made with $\mathrm{VaR}_{0.995}(R_C)$ as the target risk measure at milder target retention ratios.

The most important overarching interpretation is therefore that the tail behaviour of the severity distribution matters more for the contract choice than frequency overdispersion within the studied framework, and that the choice of risk measure is decisive for how this effect manifests itself in the treaty ranking. The study thus provides a more nuanced answer than simply stating that one contract is always best, because stop-loss is the clear winner in the chosen design at the primary target, but XL can be the more cost-effective alternative when the risk profile is dominated by extreme individual claims and the decision criterion focuses on high quantiles rather than on overall annual variation.

8 Conclusion

8.1 Answer to the research question

This thesis has examined which of proportional reinsurance, stop-loss, and per-claim excess-of-loss, from the cedent's perspective, achieves a given level of retained annual risk at the lowest expected ceded cost. Within the studied equal-risk design, where cost is measured by the expected ceded loss $\mathbb{E}[Z_C]$, the main conclusion is that stop-loss is generally the most cost-effective contract.

At the primary target retention ratio $\eta = 0.80$, stop-loss is the cheapest in all cases where a full ranking is permitted according to the 5% rule for target attainment, and the few remaining cases point in the same direction. Proportional reinsurance is usually among the most expensive contracts, while per-claim excess-of-loss generally falls between stop-loss and proportional reinsurance.

At the same time, the answer depends on the structure of the claim process. The frequency model primarily affects the level of risk and cost, but rarely changes which contract is preferable. The tail behaviour of the severity distribution is more important. For the exponential, gamma, and weibull scenarios, stop-loss is clearly the most cost-effective. In the Pareto scenarios, per-claim excess-of-loss becomes substantially more competitive and is the cheapest usable contract in certain cases when risk is measured by $\mathrm{VaR}_{0.995}(R_C)$ and the target retention ratio is milder, especially at $\eta = 0.90$ and $\eta = 0.95$, and also at $\eta = 0.85$ for the Poisson + Pareto scenario. The most robust conclusion is therefore that stop-loss is the main winner for the studied combinations of frequency-severity models and risk measures, but that per-claim excess-of-loss becomes particularly relevant when the risk profile is dominated by very large individual claims and the comparison is based on a quantile-based risk measure.

References

[1] Albrecher, H., Beirlant, J., and Teugels, J. L. (2017). *Reinsurance: Actuarial and Statistical Aspects*. John Wiley & Sons, Hoboken, NJ.

[2] Lindskog, F. (2026). *Risk Models, Claims Reserving and Solvency in Non-Life Insurance*. Lecture notes, version dated 18 January 2026.

[3] Mikosch, T. (2009). *Non-Life Insurance Mathematics: An Introduction with the Poisson Process*. Second edition. Springer-Verlag, Berlin Heidelberg.

A Additional Monte Carlo results

This appendix contains additional Monte Carlo output that supports the results in Section 6. The main text reports the most important tables and figures, while the appendix gives additional simulation settings, target-attainment diagnostics, and supplementary figures for risk measures not shown in the main results section.

A.1 Simulation settings and scenario grid

Table A.1 summarises the main simulation settings used in the final numerical experiment.

Table A.1: Main simulation settings.

Quantity	Value
Master seed	`19981118`
Training sample size	$M_{\text{train}} = 400\,000$
Test sample size	$M_{\text{test}} = 400\,000$
Confidence level	$\alpha = 0.995$
Target retention ratios	$\{0.95, 0.90, 0.85, 0.80, 0.75, 0.70\}$
Primary target retention ratio	$\eta = 0.80$
Cost proxy	$\mathbb{E}[Z_C]$
Usability rule for ranking	$\lvert\text{gap}^{\text{test}}\rvert \leq 0.05$
Tail diagnostic blocks	$B = 20$

Table A.2 gives the scenario grid used in the simulations. In all scenarios, $\mathbb{E}[N] = 10$, $\mathbb{E}[X] = 1$, and therefore $\mathbb{E}[S] = 10$.

Table A.2: Frequency–severity scenario grid.

Frequency model	Frequency parameters	Severity model	Severity parameters	Var(N)	Var(X)
Poisson	$\lambda = 10$	Exponential	$\beta = 1$	10	1.000
Poisson	$\lambda = 10$	Gamma	$a = 2,\ b = 2$	10	0.500
Poisson	$\lambda = 10$	Weibull	$\tau = 0.7,\ c \approx 1.1794$	10	2.139
Poisson	$\lambda = 10$	Pareto	$\alpha = 2.5,\ \kappa = 1.5$	10	5.000
Poisson–Gamma	$r = 5,\ \delta = 0.5$	Exponential	$\beta = 1$	30	1.000
Poisson–Gamma	$r = 5,\ \delta = 0.5$	Gamma	$a = 2,\ b = 2$	30	0.500
Poisson–Gamma	$r = 5,\ \delta = 0.5$	Weibull	$\tau = 0.7,\ c \approx 1.1794$	30	2.139
Poisson–Gamma	$r = 5,\ \delta = 0.5$	Pareto	$\alpha = 2.5,\ \kappa = 1.5$	30	5.000

A.2 Target-attainment diagnostics

The equal-risk comparison relies on the calibrated contract actually reaching the intended target risk on the independent test sample. Table A.3 summarises the target-attainment diagnostics across the full equal-risk grid.

Table A.3: Target-attainment diagnostics across the equal-risk grid.

Quantity	Value
Non-baseline equal-risk rows	576
Rows satisfying the 5% usability rule	557
Rows not satisfying the 5% usability rule	19
Contract type of all non-usable rows	XL
Sign of all non-usable test gaps	Negative
Fully rankable scenario–risk combinations at $\eta = 0.80$	28 out of 32
Non-usable rows at $\eta = 0.80$	4

Table A.4 shows the number of usable rows by target retention ratio. The calibration is most difficult for the more restrictive targets, while all rows satisfy the rule for the milder targets $\eta = 0.90$ and $\eta = 0.95$.

Table A.4: Usability by target retention ratio.

Target retention ratio η	Usable rows	Total rows
0.70	88	96
0.75	91	96
0.80	92	96
0.85	94	96
0.90	96	96
0.95	96	96

Table A.5 reports the automatic identity checks. These checks verify that the simulated retained and ceded losses satisfy the defining contract formulas.

Table A.5: Automatic numerical identity checks.

Check	Value
max $\lvert S - (R + Z)\rvert$	0
max Prop retained scaling error	0
max Prop ceded scaling error	0
max SL retained formula error	0
max SL ceded formula error	0
max SL cap violation	0
max XL retained formula error	0
max XL ceded formula error	0

A.3 Supplementary equal-risk figures

The main results section shows equal-risk curves for σ_C and $\mathrm{VaR}_{0.995}$. Figure A.1 adds the corresponding baseline curves for $\rho^{\mathrm{SD}}_{0.995}$ and $\mathrm{ES}_{0.995}$.

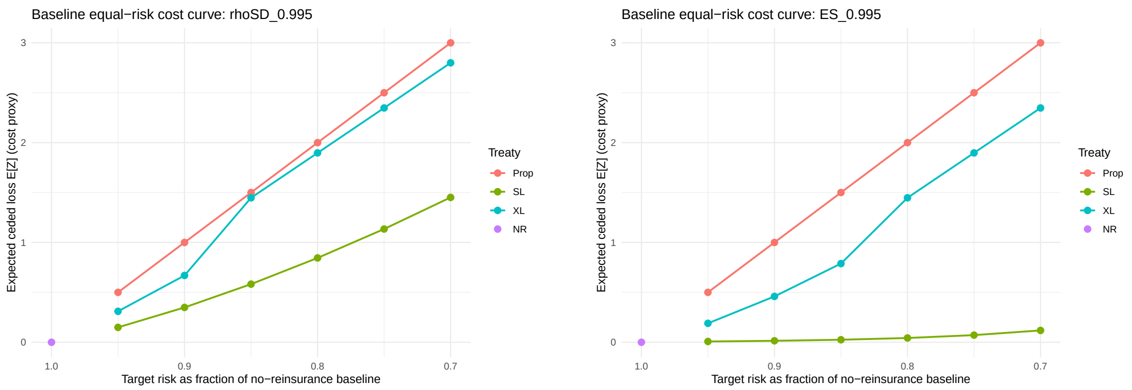

(a) $\rho^{\mathrm{SD}}_{0.995}$ (b) $\mathrm{ES}_{0.995}$

Figure A.1: Additional equal-risk cost curves for the Poisson + Exponential baseline scenario.

The main results section also shows severity comparisons under Poisson frequency for σ_C and $\mathrm{VaR}_{0.995}$. Figure A.2 adds the corresponding comparisons for $\rho^{\mathrm{SD}}_{0.995}$ and $\mathrm{ES}_{0.995}$.

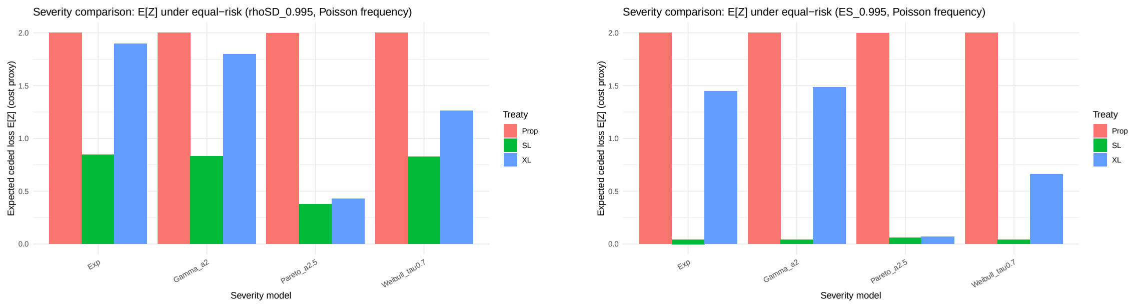

(a) $\rho^{\mathrm{SD}}_{0.995}$ (b) $\mathrm{ES}_{0.995}$

Figure A.2: Additional severity comparisons under Poisson frequency at the primary target $\eta = 0.80$.

A.4 Complementary equal-cost figures

The main analysis is the equal-risk comparison. The equal-cost comparison is used only as a complement. In this comparison, the expected ceded loss is fixed at a given cost level,

$$c = \mathbb{E}[Z_C],$$

and the resulting retained risk is compared across contracts. Since $\mathbb{E}[S] = 10$, the cost levels used in the figures are

$$c \in \{0.5, 1.0, 1.5, 2.0, 2.5, 3.0, 3.5, 4.0\},$$

corresponding to 5% to 40% of the expected gross loss. Thus, the horizontal axis in Figures A.3 and A.4 is the fixed expected ceded loss $\mathbb{E}[Z_C]$, while the vertical axis shows the retained risk measure obtained at that cost level.

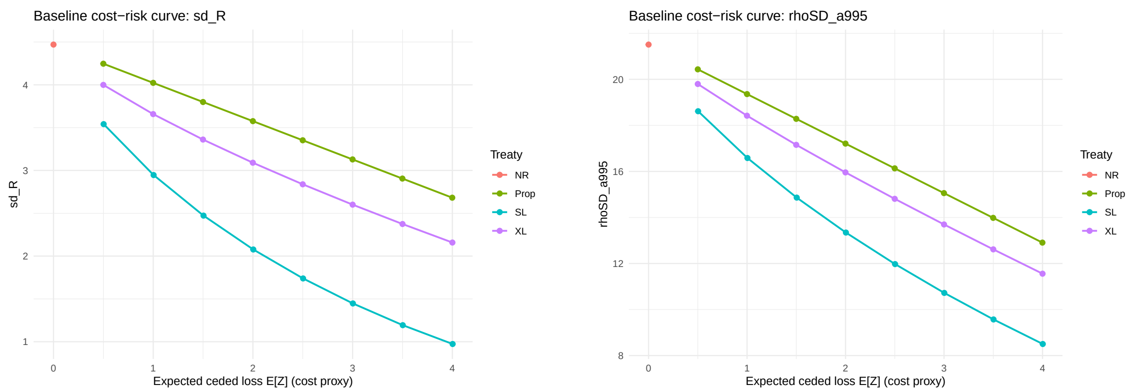

(a) σ_C (b) $\rho^{\mathrm{SD}}_{0.995}$

Figure A.3: Complementary equal-cost cost-risk curves in the Poisson + Exponential baseline scenario, part 1. The horizontal axis is the fixed expected ceded loss $\mathbb{E}[Z_C]$, used as the cost proxy.

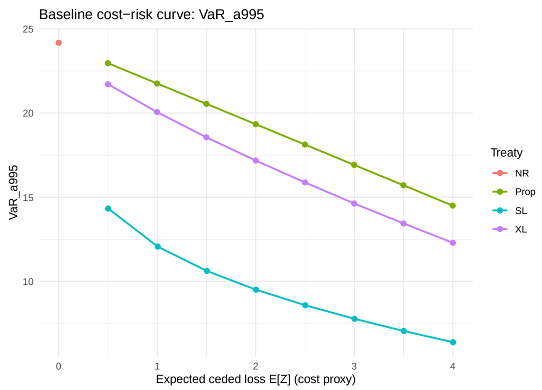

(a) $\mathrm{VaR}_{0.995}$

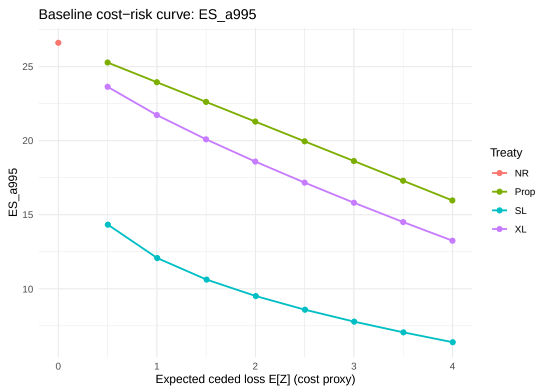

(b) $\mathrm{ES}_{0.995}$

Figure A.4: Complementary equal-cost cost-risk curves in the Poisson + Exponential baseline scenario, part 2. The horizontal axis is the fixed expected ceded loss $\mathbb{E}[Z_C]$, used as the cost proxy.