City Research Online

City, University of London Institutional Repository

Citation: Kyriakou, I. & Tsanakas, A. (2022). Efficient evaluation of alternative reinsurance strategies using control variates. *European Actuarial Journal*, 12, pp. 425-431. doi: 10.1007/s13385-022-00304-6

This is the submitted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/27414/>

Link to published version: <https://doi.org/10.1007/s13385-022-00304-6>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

City Research Online

City, University of London Institutional Repository

Citation: Kyriakou, I. and Tsanakas, A. ORCID: 0000-0003-4552-5532 Evaluating alternative reinsurance strategies using control variates. .

This is the submitted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/27266/>

Link to published version:

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Efficient evaluation of alternative reinsurance strategies using control variates

Ioannis Kyriakou and Andreas Tsanakas[†]

Faculty of Actuarial Science and Insurance, Bayes Business School (formerly Cass), City, University of London, 106 Bunhill Row, London, EC2Y 8NE, UK.

*Corresponding author(s). E-mail(s): a.tsanakas.1@city.ac.uk;

Contributing authors: ioannis.kyriakou.2@city.ac.uk;

[†]These authors contributed equally to this work.

Abstract

In this short communication, we present a new, simple control-variate Monte Carlo procedure for enhancing the evaluation accuracy of alternative reinsurance strategies that an insurance company might adopt.

Keywords: reinsurance, Monte Carlo simulation, control variates

1 Introduction

The exposure of insurance portfolios to various risk factors renders their valuation using Monte Carlo simulation a computationally intensive task. The resulting estimates can be inaccurate if the sample sizes they are based on are not sufficiently large. Increasing estimation accuracy comes with considerable increases in runtime burden, given the square-root convergence of unbiased Monte Carlo estimators. This is a notable drawback, as practicable runtimes are crucial in insurance operations, e.g., the calculation of solvency capital requirements [1], the valuation of life insurance products [2], the assessment of insurance-based investment products' risk-return profiles, the chance-risk classification of pension products [3], and the evaluation of reinsurance strategies, which represent a key risk mitigation tool for non-life insurers [4, Chapter 9]; [5].

Monte Carlo simulation convergence can be accelerated by incorporating variance reduction methods [4, Chapter 9]; [6, Chapter 3]; [2]. In this letter, we introduce a new control variate that results in a simple, yet efficient, variance reduction framework for evaluating a reinsurance strategy that a company considers to employ. The main premise of our approach is the anticipated correlation of the gross and net portfolio losses, with both variables calculated as standard outputs in model runs. This correlation is exploited to evaluate alternative reinsurance strategies efficiently, using a reduced number of expensive simulations.

2 Methodology

An insurer's portfolio loss is represented as $g(\mathbf{X}, \lambda)$, where $\mathbf{X} \sim F$ is a random vector of risk factors and λ a vector of reinsurance parameters. We will focus on three particular choices for λ :

- i) λ_0 corresponds to the case of no reinsurance, therefore $g(\mathbf{X}, \lambda_0)$ represents the *gross* portfolio loss;
- ii) λ_1 represents the current reinsurance strategy, therefore $g(\mathbf{X}, \lambda_1)$ is the *net* portfolio loss;
- iii) λ_2 represents an alternative reinsurance strategy that the company is considering, under which the *net* loss is $g(\mathbf{X}, \lambda_2)$.

The insurer is interested in quantities of the form $\mathbb{E}[\psi \circ g(\mathbf{X}, \lambda)]$, where ψ is a function that represents (tail) risk preferences, e.g., for some threshold t , $\psi(z; t) := \mathbb{1}_{\{z > t\}}$ or $\psi(z; t) := \max(z - t, 0) = (z - t)_+$, yielding, respectively, the survival function or stop-loss transform. The insurer calculates $\mathbb{E}[\psi \circ g(\mathbf{X}, \lambda_k)]$ by simulation for a given reinsurance λ_k . As part of any simulation run, the gross loss $g(\mathbf{X}, \lambda_0)$ is always evaluated by default. Hence, in each model run, the gross and net (under some λ_k) portfolio loss distribution can be worked out. We assume throughout that random number generation is computationally cheap, whilst the evaluations of g are expensive.

We distinguish between two model runs:

- *Baseline run*: we simulate i.i.d samples $\hat{\mathbf{X}}^{(i)} \sim F$, $i = 1, \dots, \hat{n}$, and evaluate $g(\hat{\mathbf{X}}^{(i)}, \lambda_1)$ and $g(\hat{\mathbf{X}}^{(i)}, \lambda_0)$.
- *Sensitivity run*: we simulate i.i.d samples $\tilde{\mathbf{X}}^{(i)} \sim F$, $i = 1, \dots, \tilde{n}$, and evaluate $g(\tilde{\mathbf{X}}^{(i)}, \lambda_2)$ and $g(\tilde{\mathbf{X}}^{(i)}, \lambda_0)$.

We require that the samples generated in the two runs are independent; achieved, for example, using a different random seed. The sample size $\hat{n}$ of the baseline run is large as we require high accuracy to calculate quantities with respect to the current portfolio, such as the firm's regulatory capital. Instead, the sample size $\tilde{n}$ of the sensitivity run is allowed to differ – for example it could be $\tilde{n} < \hat{n}$, which cuts computational time and potentially allows the investigation of multiple alternative reinsurance strategies.

Through simulation, we are interested in evaluating the change in the portfolio risk arising from moving from the current reinsurance strategy λ_1 to the alternative λ_2 . In particular, we seek to compute the quantity

$$\xi = \mathbb{E}[Y_2] - \mathbb{E}[Y_1], \quad \text{where} \\ Y_k = \psi \circ g(\mathbf{X}, \lambda_k), \quad k = 0, 1, 2.$$

To calculate ξ , we thus need to carry out both the baseline run and the sensitivity run. Our aim is to provide a variance reduction scheme that allows us to carry out the sensitivity run with sample size $\tilde{n}$ that is substantially lower than $\hat{n}$.

We define the following standard Monte Carlo estimators:

$$\hat{\mu}_0 = \frac{1}{\hat{n}} \sum_{i=1}^{\hat{n}} \hat{Y}_0^{(i)}, \quad \hat{\mu}_1 = \frac{1}{\hat{n}} \sum_{i=1}^{\hat{n}} \hat{Y}_1^{(i)}, \quad \tilde{\mu}_0 = \frac{1}{\tilde{n}} \sum_{i=1}^{\tilde{n}} \tilde{Y}_0^{(i)}, \quad \tilde{\mu}_2 = \frac{1}{\tilde{n}} \sum_{i=1}^{\tilde{n}} \tilde{Y}_2^{(i)}.$$

To estimate ξ , we define the estimator:¹

$$\check{\xi}_\beta = \tilde{\mu}_2 - \hat{\mu}_1 - \beta(\tilde{\mu}_0 - \hat{\mu}_0).$$

It is immediate that $\mathbb{E}[\check{\xi}_\beta] = \xi$. As for the variance, we have, by independence of the samples of the two runs, that

$$\mathbb{V}[\check{\xi}_\beta] = \mathbb{V}[\tilde{\mu}_2 - \beta\tilde{\mu}_0] + \mathbb{V}[\hat{\mu}_1 - \beta\hat{\mu}_0] = \frac{1}{\tilde{n}}\mathbb{V}[Y_2 - \beta Y_0] + \frac{1}{\hat{n}}\mathbb{V}[Y_1 - \beta Y_0],$$

where the equality follows by identical distribution of the samples. It then holds trivially that

$$\mathbb{V}[\check{\xi}_\beta] = \frac{\mathbb{V}[Y_2] + \beta^2\mathbb{V}[Y_0] - 2\beta\mathbb{C}[Y_2, Y_0]}{\tilde{n}} + \frac{\mathbb{V}[Y_1] + \beta^2\mathbb{V}[Y_0] - 2\beta\mathbb{C}[Y_1, Y_0]}{\hat{n}},$$

where $\mathbb{C}[\cdot, \cdot]$ denotes the covariance operator, from which the coefficient value $\beta = \beta^*$ minimising the variance of the estimator² is given by

$$\beta^* = \frac{\frac{1}{\tilde{n}}\mathbb{C}[Y_2, Y_0] + \frac{1}{\hat{n}}\mathbb{C}[Y_1, Y_0]}{\mathbb{V}[Y_0] \left(\frac{1}{\tilde{n}} + \frac{1}{\hat{n}} \right)}.$$

¹An alternative is to carry out the baseline and sensitivity runs based on the same sample size and starting from the same seed, such that the samples from the two runs are not independent, that is, estimate ξ by $\hat{\mu}_2 - \hat{\mu}_1$. Such use of common samples [7, Chapter 8], has two potential drawbacks. First, as $\hat{\mu}_2$ is also estimated on a sample of size $\hat{n}$, there are no savings in computational time. Second, the use of common numbers to evaluate $\hat{\mu}_2$ and $\hat{\mu}_1$ may be impractical, depending on the commercial software platform used, as a change from λ_1 to λ_2 could affect the simulated sequences, even if the same seed is used (e.g., if g is itself evaluated numerically).

²In practice, the optimal β^* value is estimated via its sample counterpart, which is likely to introduce some bias, still usually small [8]. A simple way to eliminate this bias is by estimating β^* independently using a pilot sample – we do not pursue this route further here.

Finally, we get

$$\mathbb{V}[\check{\xi}_{\beta^*}] = \frac{\mathbb{V}[Y_2]}{\tilde{n}} + \frac{\mathbb{V}[Y_1]}{\hat{n}} - \frac{\left(\frac{1}{\tilde{n}}\mathbb{C}[Y_2, Y_0] + \frac{1}{\hat{n}}\mathbb{C}[Y_1, Y_0]\right)^2}{\mathbb{V}[Y_0] \left(\frac{1}{\tilde{n}} + \frac{1}{\hat{n}}\right)}.$$

Note that in the special case of the naive estimator, we have $\check{\xi}_0 = \tilde{\mu}_2 - \hat{\mu}_1$ with $\mathbb{E}[\check{\xi}_0] = \xi$ and $\mathbb{V}[\check{\xi}_0] = \mathbb{V}[Y_2]/\tilde{n} + \mathbb{V}[Y_1]/\hat{n}$.

If the gross and net portfolio loss have a high correlation, as is typical in practice, we can achieve a substantial reduction in the variance of our estimator with respect to the naive one. It is helpful to have a simple rule of thumb which would indicate the required correlation level for a modest size $\tilde{n}$ to yield sufficient reduction in variance and, thus, potential computational savings. Upon requiring the variance estimate of $\check{\xi}_{\beta^*}$ to be consistent with that of $\hat{\mu}_1$, an approximate calculation, not documented here, yields

$$\frac{\tilde{n}}{\hat{n}} \approx \frac{1}{\rho[Y_1, Y_0]^2} - 1,$$

where $\rho[Y_1, Y_0] := \mathbb{C}[Y_1, Y_0]/\sqrt{\mathbb{V}[Y_0]\mathbb{V}[Y_1]}$. Hence, it follows that $\tilde{n} < \hat{n}$ when $|\rho[Y_1, Y_0]| > \sqrt{2}/2$. This is the same benchmark correlation value that appears in the standard control-variate setting [7, Chapter 8].

3 Example

Here, we illustrate the use of the estimator $\check{\xi}_{\beta^*}$ in a toy example. Consider a homogeneous portfolio with 10 lines of business, $\mathbf{X} = (X_1, \dots, X_{10})$, and portfolio loss

$$g(\mathbf{X}, \lambda) = \sum_{j=1}^{10} (X_j - R_j + \mathbb{E}[R_j] + 0.5\sigma[R_j]),$$

where:

- X_j are gross losses for each line of business, which we assume to be identically log-normally distributed, with $\mathbb{E}[X_j] = 100$, $\sigma[X_j] := \mathbb{V}[X_j]^{1/2} = 20$, and dependent via a Gaussian copula with correlation parameter 0.3 for all pairs.
- $R_j = \min((X_j - d)_+, l)$ is the reinsurance recovery for each line of business and $\mathbb{E}[R_j] + 0.5\sigma[R_j]$ is the reinsurance premium. We, thus, have that $\lambda = (d, l)$. Specifically, for the gross loss we have $\lambda_0 = (0, 0)$, for the net loss of the baseline run we have $\lambda_1 = (F_{X_j}^{-1}(0.7), F_{X_j}^{-1}(0.8) - F_{X_j}^{-1}(0.7))$, whereas for the sensitivity run we have $\lambda_2 = (F_{X_j}^{-1}(0.75), F_{X_j}^{-1}(0.85) - F_{X_j}^{-1}(0.75))$. Hence, the sensitivity run aims at evaluating the effect of making all reinsurance layers somewhat higher.

We simulate from this model, with baseline run sample size $\hat{n} = 10^5$ and sensitivity run sample size $\tilde{n} = 10^4$, and derive estimates of the quantities $\mathbb{E}[Y_k]$, $k = 0, 1, 2$, where

$$\begin{aligned} Y_k &= \psi(g(\mathbf{X}, \lambda_k); t_k) \\ \psi(z; t_k) &= \begin{cases} \mathbb{1}_{\{z > t_k\}} & (\text{survival function}); \text{ or} \\ (z - t_k)_+ & (\text{stop-loss transform}) \end{cases} \\ t_0 &= F_{Y_0}^{-1}(\alpha), \quad t_1 = t_2 = F_{Y_1}^{-1}(\alpha), \quad \alpha \in [0.9, 0.995]. \end{aligned}$$

We compare the performance of the naive Monte Carlo estimator $\check{\xi}_0$ with the improved one $\check{\xi}_{\beta^*}$, by evaluating both on $m = 500$ blocks of simulated samples. The results of this exercise are summarised in Figure 1, for different thresholds t_2 , corresponding to portfolio net loss levels. The top plots (a, b) deal with estimating the survival function under the alternative reinsurance strategy λ_2 , whereas the bottom plots (c, d) show the estimation of the stop-loss transform. On the left (a, c), we show realisations of the difference estimators $\check{\xi}_0$ (in grey) and $\check{\xi}_{\beta^*}$ (in blue). On the right, we show estimates of the survival function (b) and the stop-loss transform (d) under the reinsurance strategy λ_2 , together with 95% confidence intervals evaluated across the $m = 500$ experiments. From all plots, it is seen that the control-variate estimator $\check{\xi}_{\beta^*}$ exhibits a major improvement compared to the naive estimator $\check{\xi}_0$ – more so in the case of the stop-loss transform. From experiments not documented here, we found that the same holds after changes in the number of lines of business, their volatility, and their copula/correlation specification.

Nonetheless, one part of the problem specification that does impact performance is the reinsurance strategy λ_k , $k = 1, 2$; if these λ_k are such that the correlations (Y_0, Y_k) are not high enough, adopting the control-variate estimator $\check{\xi}_{\beta^*}$ may not offer a substantial benefit. For example, in the simulation study above, for survival function estimation and threshold loss levels derived with $\alpha = 0.95$, we have $t_1 = t_2 = 1197$, correlations $\rho[Y_1, Y_0] = 0.956$, $\rho[Y_2, Y_0] = 0.948$ and variance reduction factor $\sigma[\check{\xi}_{\beta^*}]/\sigma[\check{\xi}_0] = 0.31$. Alternatively, let us consider the case of $\lambda_2 = (F_{X_j}^{-1}(0.6), F_{X_j}^{-1}(0.85) - F_{X_j}^{-1}(0.6))$, such that much more of the company's risk is ceded to reinsurers. Then, for the same loss thresholds, we have $\rho[Y_2, Y_0] = 0.796$ and $\sigma[\check{\xi}_{\beta^*}]/\sigma[\check{\xi}_0] = 0.57$, reflecting a more modest variance reduction.

4 Conclusion

We presented a simple approach for improving the accuracy of simulation-based risk estimates when different model runs, as a default, produce samples from two correlated outputs (in this case, gross and net loss). By reducing computational costs, insurers can expand the scope of their exploration of

Fig. 1 Performance of naive and improved estimators against net portfolio losses (at thresholds t_2). (a, b): estimation of the portfolio survival function; (c, d): estimation of the portfolio stop-loss transform; (a, c): realisations of $\check{\xi}_0$ and $\check{\xi}_{\beta^*}$; (b, d): risk estimates with 95% confidence intervals (log-log scale).

alternative reinsurance strategies, resulting in more efficient portfolio structures. The proposed method is easily applicable in practice, as it is generic and not tailored to any particular stochastic model specification.

Statements and Declarations

Competing Interests and Funding: none.

References

- [1] Krah, A.-S., Nikolić, Z., Korn, R.: A least-squares Monte Carlo framework in proxy modeling of life insurance companies. *Risks* **6**(2), 1–26 (2018)
- [2] Korn, R., Temoçin, B.Z., Wenzel, J.: Applications of the central limit theorem for pricing Cliquet-style options. *European Actuarial Journal* **7**(2), 465–480 (2017)
- [3] Graf, S., Korn, R.: A guide to Monte Carlo simulation concepts for assessment of risk-return profiles for regulatory purposes. *European Actuarial*

Journal **10**(2), 273–293 (2020)

- [4] Albrecher, H., Beirlant, J., Teugels, J.L.: Reinsurance: Actuarial and Statistical Aspects. Wiley Series in Probability and Statistics. John Wiley & Sons, Hoboken (2017)
- [5] Zanolto, A., Clemente, G.P.: An optimal reinsurance simulation model for non-life insurance in the Solvency II framework. *European Actuarial Journal* (2021)
- [6] Korn, R., Korn, E., Kroisandt, G.: Monte Carlo Methods and Models in Finance and Insurance. Chapman and Hall/CRC Financial Mathematics Series. CRC Press, Boca Raton (2010)
- [7] Owen, A.B.: Monte Carlo Theory, Methods and Examples, (2013). Available online at <https://artowen.su.domains/mc/Ch-var-basic.pdf>
- [8] Glasserman, P.: Monte Carlo Methods in Financial Engineering. Stochastic Modelling and Applied Probability. Springer, New York (2004)