December 2013

Using leakage as a control variate

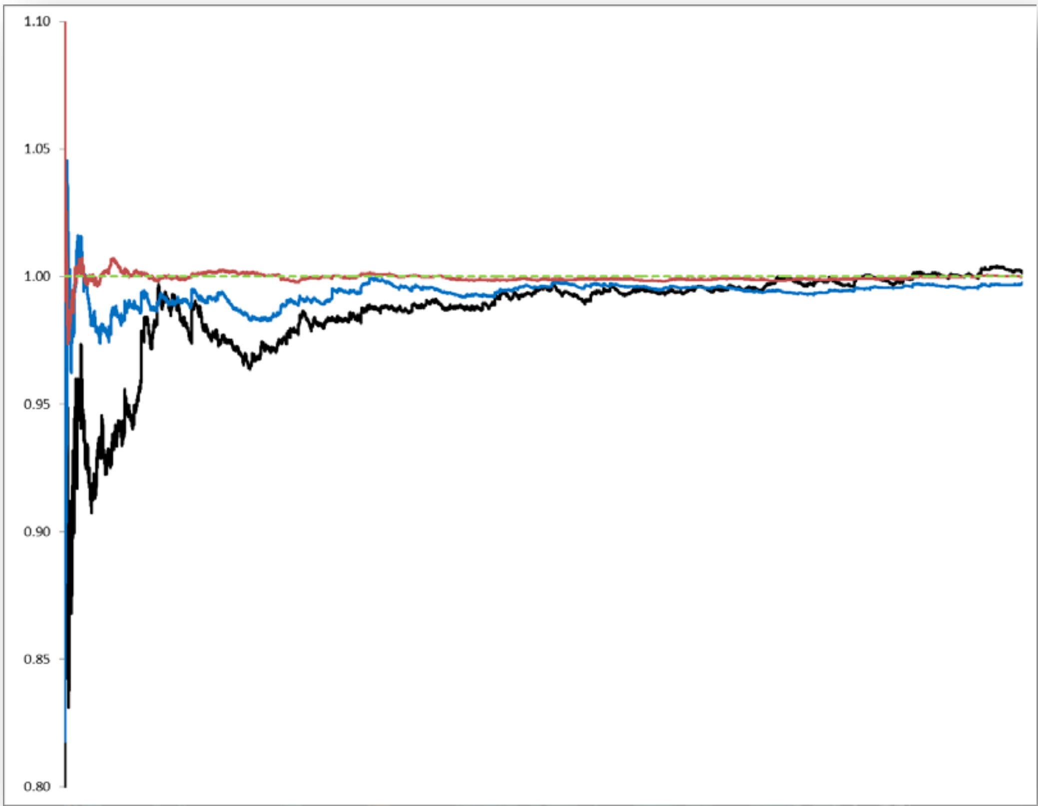

Dr. Andreas Meister, Dr. Marcel Ambrus, Dr. Urs Burri – Baloise Switzerland (Basel)

Dr. Mario Hörig, Tigran Kalberer - Milliman (Düsseldorf/Zürich)

Table of Contents

TABLE OF CONTENTS 2

1 EXECUTIVE SUMMARY 3

1.1 Background and objectives 3

1.2 Scope and structure of this report 3

1.3 Main findings 3

2 INTRODUCTION 4

2.1 Motivation – Using control variates as a tool for variance reduction 4

2.2 Leakage – Definition and reasons 5

3 USING THE LEAKAGE AS A CONTROL VARIATE 6

3.1 Setup of the case study 6

3.2 Case 1 – Without antithetic pairs 7

3.2 Case 2 – With antithetic pairs 9

4 REFERENCES AND CONTACT DETAILS 12

4.1 References 12

4.2 Contact details 12

1 EXECUTIVE SUMMARY

1.1 Background and objectives

Variance reduction is a vital part of the valuation of insurance liabilities within a Monte Carlo simulation framework. The complex nature of insurance products and their exposure towards a variety of risk factors—which usually are stochastically simulated—typically leads to a large number of scenarios, required to derive estimates for the expected value of these liabilities. The reliability of these estimates is measured by a confidence interval, normally in terms of the standard deviation of the sampled values. Well-known square root convergence, however, implies that reducing the confidence interval of the estimates by increasing the number of simulations is not very efficient because using N times the number of simulations will reduce the confidence interval only by a factor of $\frac{1}{\sqrt{N}}$. This is a serious drawback since the run-times involved in the valuation of the insurance liabilities are a crucial factor for the steering and management of insurance companies. Therefore, insurers turn towards variance reduction techniques that allow for scaling down the remaining variance without increasing the overall simulation budget. While the implementation of these methods typically requires a change of the overall simulation framework, we present in this document a new control variate approach that works within the existing simulation framework. It achieves a reduction of variance by combining the leakage of the simulation model with the liability value under consideration in a smart way.

1.2 Scope and structure of this report

Section 2 serves as an introduction to the concepts of variance reduction and leakage and aims at giving some background for the subsequent case study in Section 3.

Section 2.1 motivates the use of variance reduction methods. It briefly describes popular techniques and focuses on the control variate approach. Section 2.2 deals with the definition of the leakage in valuation simulations for insurance liabilities and describes typical reasons for its occurrence.

Section 3 introduces the main idea of this paper, i.e., to use leakage as a control variate, and illustrates the power of this technique with two case studies in Sections 3.2 and 3.3. The case studies were carried out using a real portfolio of life insurance liabilities in an actuarial projection model that is used for different operational tasks such as MCEV, Swiss Solvency Test, IFRS and ALM.

Section 3.1 describes the setup of the case studies.

1.3 Main findings

This report demonstrates that model leakage can be efficiently used as a control variate to reduce the variance of economic valuations of insurance liabilities if this model leakage displays a certain level of correlation to the liability value under consideration. We also show that the control variate approach utilizing leakage can be applied on top of existing variance reduction techniques such as antithetic sampling of the economic scenarios. Again, the level of variance reduction achieved depends on the level of correlation between leakage and liability value.

The application of our technique is straightforward and generic because:

- The approach consists of a simple algebraic adjustment of the liability values in each scenario.
- All ingredients for this adjustment are already available within the simulation and reporting framework.

2 INTRODUCTION

2.1 Motivation – Using control variates as a tool for variance reduction

Following Glasserman [GL], variance reduction techniques can be classified into two different categories:

(i) Techniques trying to reduce the variability of simulation input
(ii) Techniques taking advantage of tractable features of a model to adjust or correct simulation outputs

Examples for (i) are antithetic or stratified sampling approaches while the control variate approach is a classic example for (ii). The main advantage of approaches of category (i) is that many of them can be applied in a generic way without paying attention to the particular nature of the simulation. Category (ii) methods cannot be applied in such a generic way because they need some specific knowledge of the variables and the simulation scheme under consideration. If they are available, though, they are often much more powerful than the techniques from category (i) and lead to a substantial reduction of variance as they are well suited to the particular situation.

Now we turn to the theoretical foundations of using a control variate for variance reduction and follow [GL] in the exposition:

Let Y be a random variable with mean μ_Y. Given n realizations $y_1,\ldots,y_n$ of Y we start by defining

$$\widehat{\mu_Y} = \frac{1}{n}\sum_{i=1}^{n} y_i \qquad (1)$$

as an unbiased estimate for μ_Y. Now suppose that there is a random variable X with known mean μ_X and a non-zero correlation[1] to Y. Let $x_1,\ldots,x_n$ be realizations of X that have been drawn jointly with Y, i.e. (x_i,y_i) is a (joint) realization of the random vector (X,Y) and are based on the same random input. We now want to take advantage of the correlation between X and Y to correct the single realizations y_i and adjust them via

$$\tilde{y}_i = y_i - b(x_i - \mu_X). \qquad (2)$$

for a scalar b regarded as optimal (see below).

The idea behind (2) is to measure the deviation between x_i and its mean value and subtract this difference (multiplied with a scalar b) from y_i, since once this deviation is large the correlation between X and Y implies that we have a certain likelihood that the deviation between y_i and μ_Y is large as well. Therefore, we measure the 'noise' per simulation, i.e., the difference between X and its mean, and try to use it to correct the realizations of Y by a multiple of this quantity utilizing the correlation between X and Y. Furthermore, note that $\widehat{\mu_{\tilde{Y}(b)}} = \frac{1}{n}\sum_{i=1}^{n} \tilde{y}_i$ is an unbiased estimator for μ_Y (follows from (2)).

The remaining question is how to choose the scalar b in an optimal way. 'Optimal' in this sense means that the resulting variance of $\widehat{\mu_{\tilde{Y}(b)}} = \frac{1}{n}\sum_{i=1}^{n} \tilde{y}_i = \widehat{\mu_Y} - b\left(\frac{1}{n}\sum_{i=1}^{n} x_i - \mu_X\right)$ is minimal. The variance of $\widehat{\mu_{\tilde{Y}(b)}}$ is

$$\mathrm{Var}\left(\widehat{\mu_{\tilde{Y}(b)}}\right) = \mathrm{Var}(\widehat{\mu_Y}) + \frac{1}{n}\left(b^2\mathrm{Var}(X) - 2b * \mathrm{Cov}(X,Y)\right) \qquad (3)$$

and is hence minimized by

$$b^* = \frac{\mathrm{Cov}(X,Y)}{\mathrm{Var}(X)}. \qquad (4)$$

We obtain the variance

[1] Note that the sign of the correlation is irrelevant for this application.

$$\mathrm{Var}\left(\widehat{\mu_{\tilde{Y}(b^*)}}\right) = (1 - \mathrm{Cor}(X,Y)^2) * \mathrm{Var}(\widehat{\mu_Y}) \ . \quad (5)$$

Analyzing (5) we see that the resulting variance scales linearly with the original variance of $\widehat{\mu_Y}$ and decreases with increasing correlation between X and Y. The decrease of the variance achieved by the control variate thus depends on the correlation between X and Y where the quadratic relationship in (5) indicates that the strength of this effect strongly increases once the correlation approaches either 1 or -1.

Note that equation (4) has an intuitive geometric interpretation: it is the slope of the regression line fitting Y to X.

Remark: Typically, both ingredients for the construction of b^*, Cov(X,Y) and Var(X), are calculated on the sample $\{(x_i,y_i), i=1,\dots,n\}$ and hence need not be known in advance.

In the context of the indirect valuation of insurance liabilities, Y typically refers to the 'portfolio value,' e.g., the net present value of future shareholder profits (NPVFP); in this case, a popular choice for a control variate is the (pathwise) value of a replicating portfolio,[2] calibrated to the present value of the shareholder profits. By definition, a replicating portfolio will display a high correlation to the liability value under consideration and hence achieve a high level of variance reduction via (5). However, the application of this method requires the existence of a robust and reliable replicating portfolio which itself is accompanied by certain simulation and calibration efforts. Therefore, the application of the control variate approach so far has been restricted to insurers using a replicating portfolio. As opposed to this, the control variate approach presented in this paper is always applicable in the setting of valuation of liabilities. The extent of the variance reduction depends on the correlation between the NPVFP and the proposed control variate.

2.2 Leakage – Definition and reasons

The (market-consistent) valuation of insurance liabilities is based on the simultaneous projection of assets and liabilities as well as their interaction, where the asset performances follow the risk-neutral measure (see [KA] for details). Under the risk-neutral measure, the discounted price process of any such asset is a martingale. This central property of risk-neutral valuation for assets implies that the initial market value of the assets must in theory coincide with the expected value of the discounted total cash flows arising from the projected liabilities[3] plus the discounted value of the market value of the residual assets at the end of the projection.

Any difference between these two quantities is called leakage and can be regarded as an indicator of the 'valuation error' within the economic liability projection.

The reasons for the occurrence of leakage can typically be classified as follows:

(i) Sampling error due to using a finite number of scenarios.
(ii) Arbitrage opportunities within the economic scenario generator (ESG), i.e., assets projected within the ESG not showing the crucial martingale property mentioned above.
(iii) Inconsistencies within the cash flow model. Examples are:
- Cash flows not being attributed to one of the three stakeholders (policyholder, shareholder, tax) and hence leaking from the projection model.
- Timing issues: e.g., suppose a EUR cash flow occurring in the course of the year is converted to CHF using the end-year exchange rate and then rolled-up to end-year using the CHF cash return. This will result in an inconsistency if CHF and EUR cash returns are different.
- If market values of some assets are calculated by using analytic formulae based on a different economic model than the one underlying the scenario generator.

[2] To be more precise, the present value of the cash flows arising from the assets of the replicating portfolio.

[3] These are typically cash flows to/from the policyholder such as premia and benefits, cash flows from/to the shareholder such as capital injections and dividends, and tax payments. For the sake of simplicity, expenses are considered as policyholder cash flows in this paper.

While reasons (ii) and (iii) can be avoided in principle, reason (i) is an intrinsic property of Monte Carlo simulations and can only be addressed by reducing the variance of the simulation scheme. For the purpose of this paper it is fundamental that a significant part of the leakage within a cash flow projection can be traced back to reason (i).

3 USING THE LEAKAGE AS A CONTROL VARIATE

Having introduced control variates as a suitable variance reduction technique and the concept of leakage within the valuation of insurance liabilities, we will now combine these two in order to achieve a simple but efficient variance reduction.

At the heart of this section is the observation that leakage and liability values (such as the net present value of future shareholder profits) often display a high correlation on a set of risk-neutral valuation scenarios (see, e.g., Figure 1 below). The subsequent case study illustrates this finding and its applicability by constructing a leakage-based control variate that allows for a reduction of the variance in the estimation of the expected value of the liabilities.

While we might find some intuitive explanation for the fact that leakage and liability values are correlated (the pricing error scales similarly to the value of the liabilities), we want to emphasize that a high correlation is by no means automatically granted. Smaller correlations than in our case study may be encountered if other capital market models and calibrations are used or if other liabilities are projected.

3.1 Setup of the case study

The case study illustrates the power of using leakage as a control variate by applying this technique to a real portfolio of liabilities, modeled within a realistic cash-flow model using stochastic economic scenarios.

The portfolio of liabilities is a Swiss group pension scheme comprising of:

- Annuities (retirees, widows, disabled)
- A mandatory and a non-mandatory saving process
- Optional conversion to annuities of the savings on reaching retirement
- Risk and expense process

The size of the portfolio is roughly given by about 1 bio CHF annual premium and 1 bio CHF single premium.

The main dynamic features are:

- Dynamic policyholder behavior
    - Contract lapse with guaranteed redemption value
    - Capital take-up rate at reaching retirement
- Regulatory requirements on the policyholder participation via a legal quote involving all processes: strengthening of annuity reserves, the saving, risk and expense processes
- Dynamic projection of the guaranteed interest rates for the saving processes and the conversion factors for the conversion of savings to annuities at retirement

The modeled asset portfolio (totaling a market value of about 15 bio CHF) contains assets in two economies (CHF and EUR) and consists of about 80% of bonds and mortgages, about 15% of property and about 5% of other assets like equities, hedge funds, etc.

The portfolio value under consideration is the NPVFP, the net present value of future shareholder profits. In order to render the values unrecognizable, the expected NPVFP was shifted to 1 bio CHF.

For the projection of the assets we use a two-factor Libor market model for nominal interest rates and geometric Brownian motions for the other projected indices. The overall capital market model is calibrated to market conditions as of December 31, 2012.

Part 2 of the case study uses the ordinary cash account, while Part 1 uses the zero-coupon bond with maturity 40 years as numeraire. The reason for this is that the change of numeraire leads to a significantly higher correlation between leakage and the NPVFP, which can be exploited by our control variate technique. For Part 1, we start with 25,000 stochastic scenarios where no further variance reduction technique has been applied, illustrate the correlation between leakage and NPVFP, construct a minimal-variance leakage-based control variate, apply it to correct the NPVFP per scenario and show how this helps to decrease the resulting variance in estimating the mean value of the NPVFP. The second case study starts with 25,000 antithetic pairs of stochastic scenarios. There, we first use the variance reduction provided by the antithetic pairs and then apply the leakage-based control variate on top of it to demonstrate that these two methods can be very well combined.

3.2 Case 1 – Without antithetic pairs

We start with 25,000 valuation scenarios without further intrinsically applied variance reduction. If N_j and L_j denote the contribution of scenario j to NPVFP and leakage, respectively, then the adjusted (i.e., leakage-corrected) contribution to the NPVFP is given by $N_j - b^* L_j$. Note that:

- L_j equals the initial market value of the assets minus the present value of the discounted total cash flows arising from the projected liabilities in scenario j minus the discounted value of the market value of the residual assets at the end of the projection of scenario j.
- in this definition we have made use of our assumption that the expected value of the leakage is zero.[4] This means that reasons (ii) and (iii) for the occurrence of leakage can be ruled out in the model under consideration, a claim that has to be supported by evidence from other sources than the sample considered. In the case of the portfolio of liabilities used in this case study, we are confident that reason (iii) does not give rise to any significant leakage: The model has been in use for a number of years with a variety of economic scenario generators. The leakage has been closely monitored from the start, and any indications to inconsistencies of type (iii) have been followed up and the inconsistencies have been removed. As for reason (ii), we are relying on the quality of a widely used economic scenario generator of an external provider. The scenarios have also been validated by passing a number of martingale tests.

Figure 1 displays a scatter plot of the leakage L_j and the NPVFP contributions N_j in the 25,000 scenarios (blue points), clearly showing a strong correlation. Here, the correlation between the two quantities is 96.1%, and the optimal scaling factor b^* as derived by (4) equals 0.116 and coincides with the slope of the regression line (in orange) through the 25,000 points. The red points display the adjusted contributions $N_j - b^* L_j$ to the NPVFP. With the interpretation of b^* as the slope of the regression line fitting N_j to L_j, the expression $N_j - b^* L_j$ removes the L_j-dependence from the N_j (in the sense that the adjusted values have correlation zero to the leakage), reducing the variability of those values and hence the variance. In order to visualize this reduction, the 1% / 99% percentiles of the two sets of contributions to the NPVFP are shown in the diagram.

The standard deviation of the mean value estimate is 9.35 mio CHF in the unadjusted and 2.58 mio CHF in the control variate case, and hence we have an improvement from using the control variate by a factor of 3.62. A similar factor can be observed in the narrowing of the percentile band. Note that without applying the control variate or other variance reduction techniques, such an improvement could only be achieved by using 13 times the number of valuation scenarios, i.e., in our case about 325,000 scenarios.

[4] If the expected value of the leakage differs from zero but can be specified, it should be taken into account in the adjusted NPVFP according to formula (2). Such a situation can occur in the case of an inconsistency of type (iii), third bullet point in section 2.2.

Figure 1: Scatter Plot Leakage vs. NPVFP – Case 1

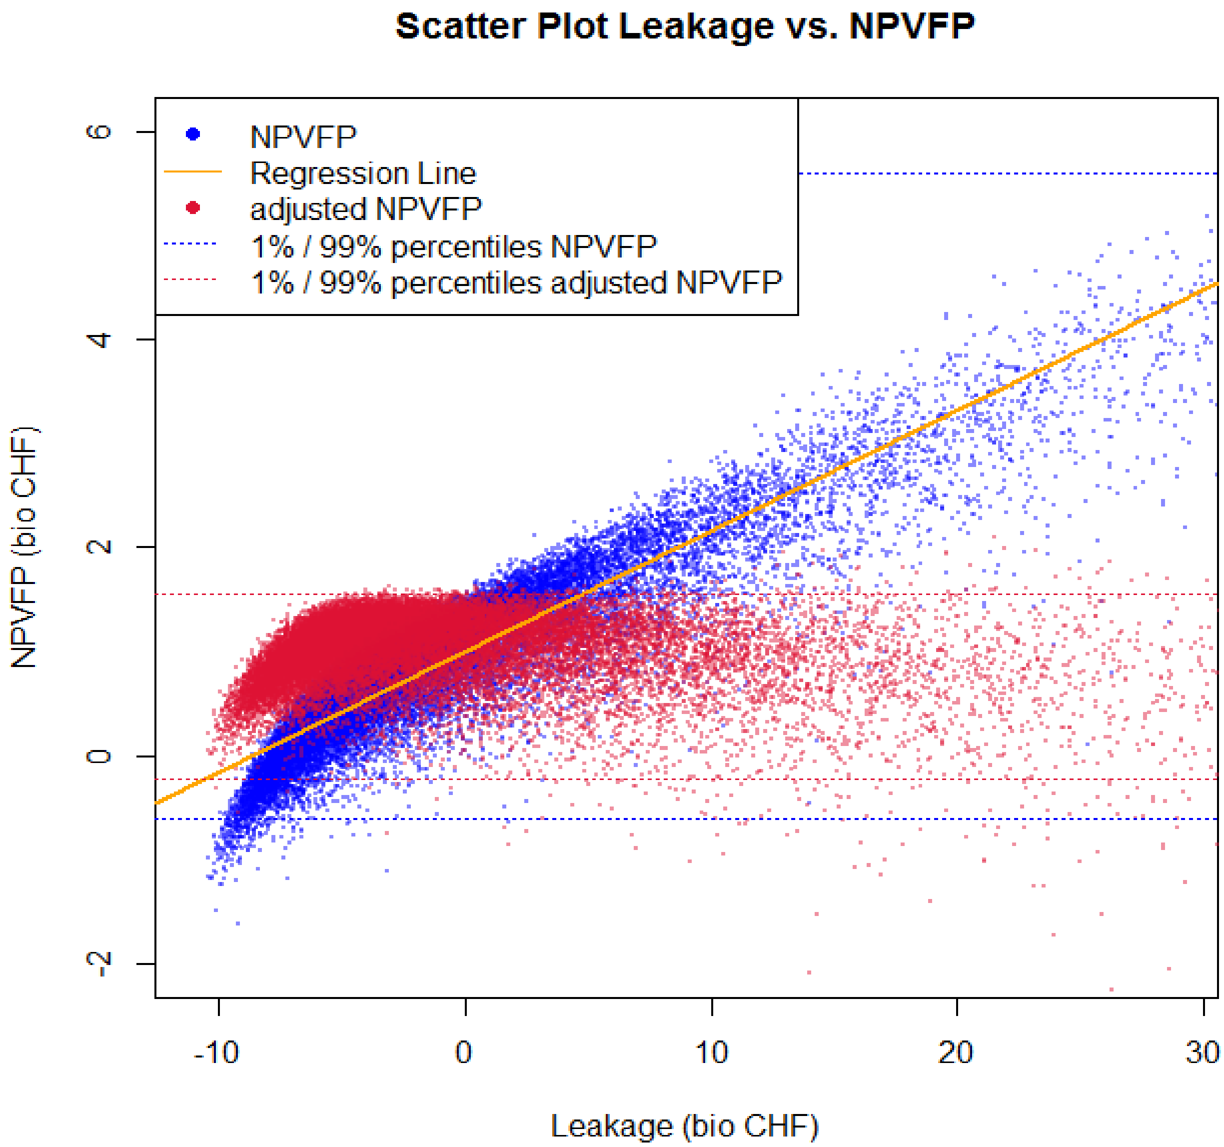

The effect of applying the leakage as a control variate can be further illustrated by convergence plots. They show the two quantities

- $\mathrm{NPVFP}\ (n) = \frac{1}{n}\sum_{j=1}^{n} N_j$
- $\mathrm{NPVFP}_{adj}(n) = \frac{1}{n}\sum_{j=1}^{n} (N_j - b^* L_j)$

as n runs from 1 to 25,000.

Figure 2: Convergence Plot With and Without Leakage-based Control Variate – Case 1

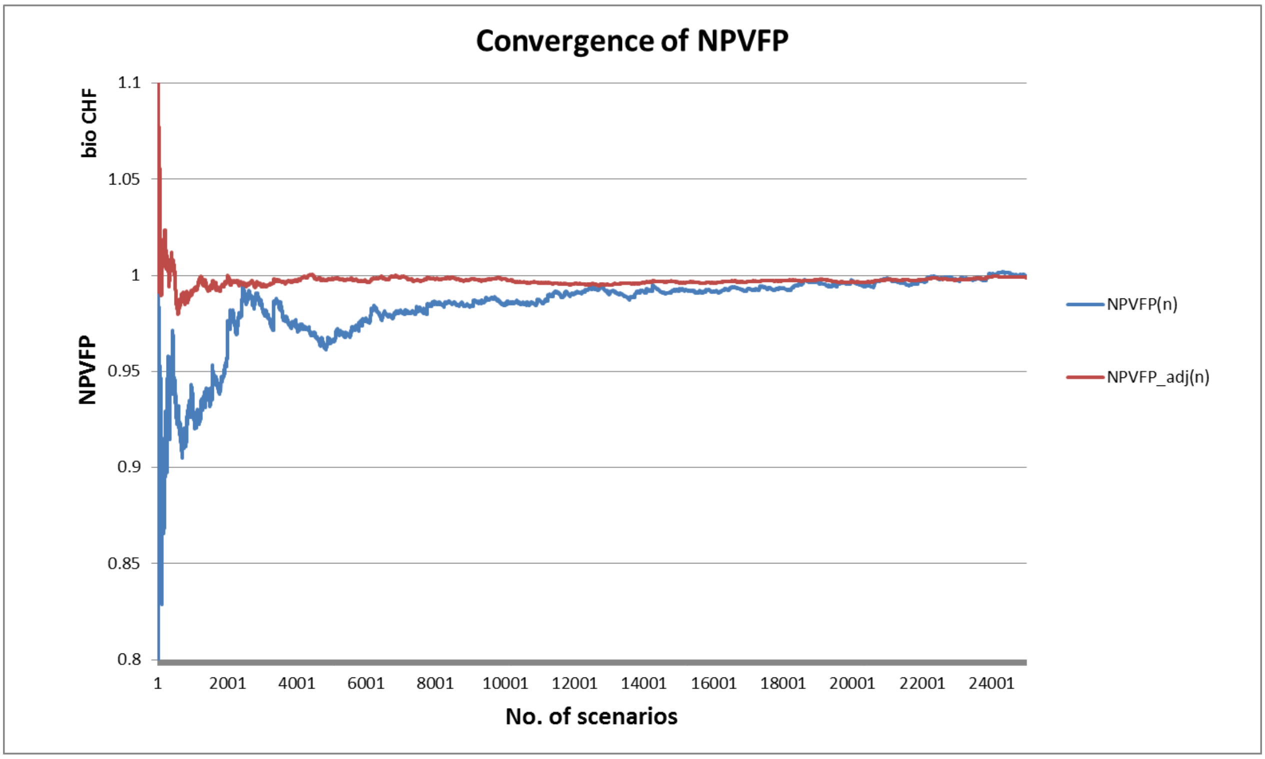

The convergence plot in figure 2 illustrates the significantly increased speed of convergence of the adjusted estimate for the mean NPVFP. While we clearly see that both estimates converge to about the same level of 1 bio CHF, it is evident that, using the original data, we basically need all the 25,000 simulations to achieve a sufficient degree of convergence, whereas the data adjusted according to the control variate reaches a stable level of convergence about four to five times faster, i.e., after only 5,000 to 6,000 scenarios.

3.2 Case 2 – With antithetic pairs

For the second case study, 50,000 scenarios consisting of 25,000 antithetic pairs have been used. To demonstrate how the leakage-based control variate can be applied on top of the variance reduction due to the antithetic pairs, we start by taking the average over each of the antithetic pairs. This results in 25,000 realizations of contributions to the NPVFP and leakage. In a second step, we make use of the leakage as a control variate on these 25,000 realizations. An interesting point about this case is that, contrary to Case 1, on the full set of 50,000 scenarios the correlation between NPVFP and leakage is very small (about 8%), so applying the control variate directly does not lead to any significant improvement.

Figure 3 displays a scatter plot of the leakage and NPVFP values of these 25,000 new realizations. Again, we see a clear correlation between leakage and NPVFP, which is however less pronounced than for Case 1. Here, the correlation between the two quantities is 69.8%, and the optimal scaling factor b^* as derived by (4) equals 1.168. Again, the orange line in Figure 3 is the regression line through the 25,000 and the red points display the adjusted contributions to the NPVFP.

Figure 3: Scatter Plot Leakage vs. NPVFP – Case 2

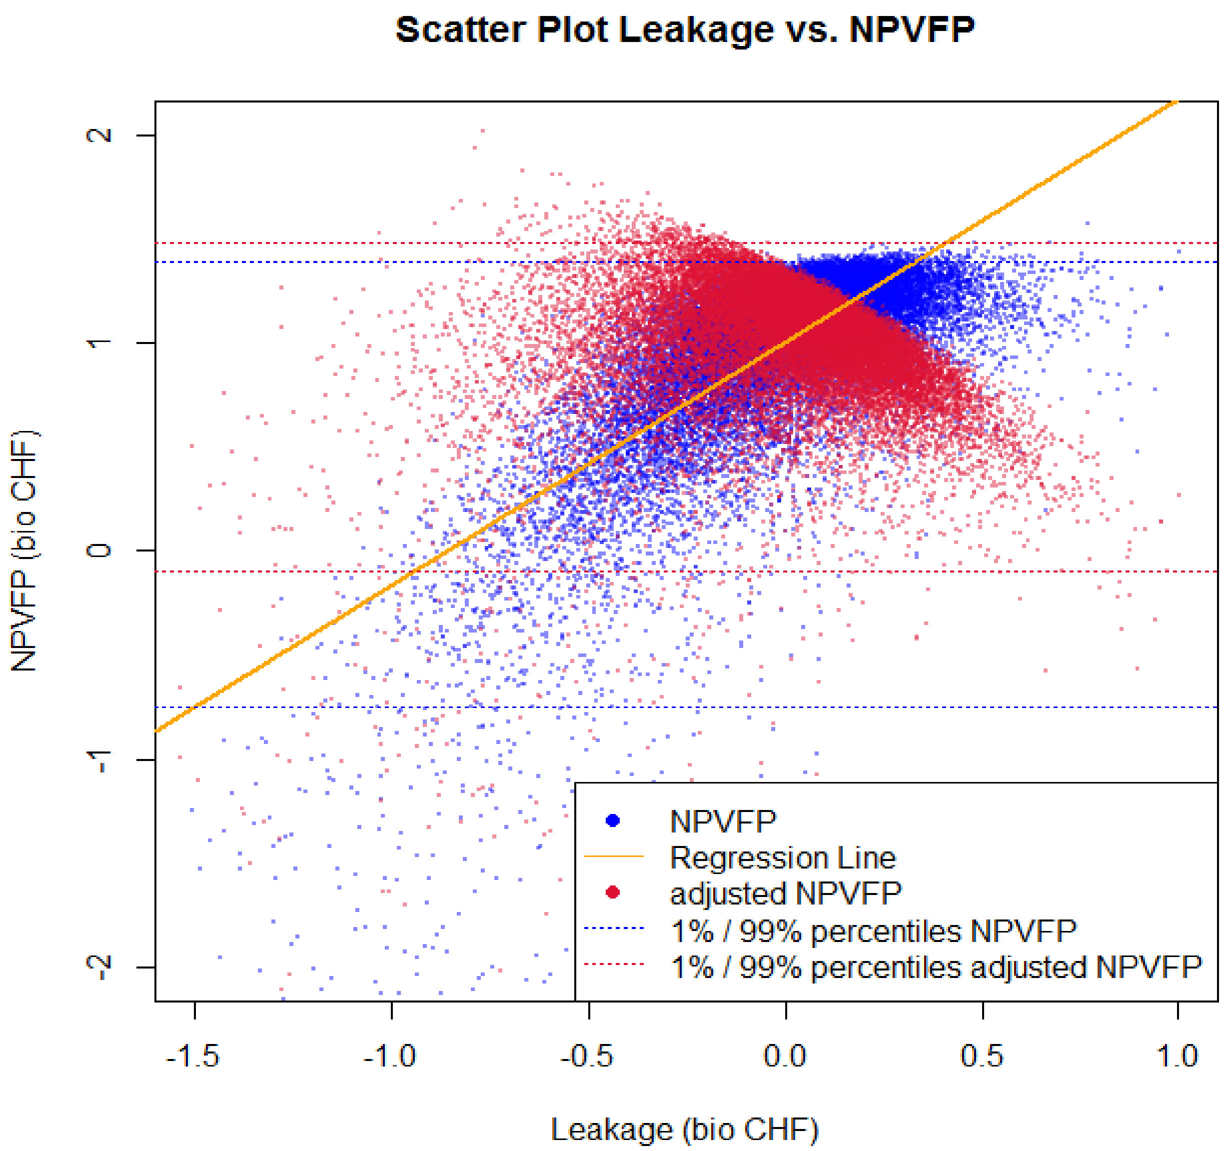

The standard deviation of the mean value estimator is 2.81 mio CHF in the unadjusted antithetic and 2.01 mio CHF in the combined antithetic and control variate case. Again, this reduction by about 0.7 can also be seen in the narrowing of the percentile band. The variance reduction achieved by using the leakage as a control variate on top of the antithetic scenarios is smaller than in Section 3.1 but is still significant. It reduces the necessary simulation budget for the same accuracy by a factor of about two.

Now we turn to the convergence plot.

Figure 4: Convergence Plot for Antithetic Pairs With and Without Leakage-based Control Variate – Case 2

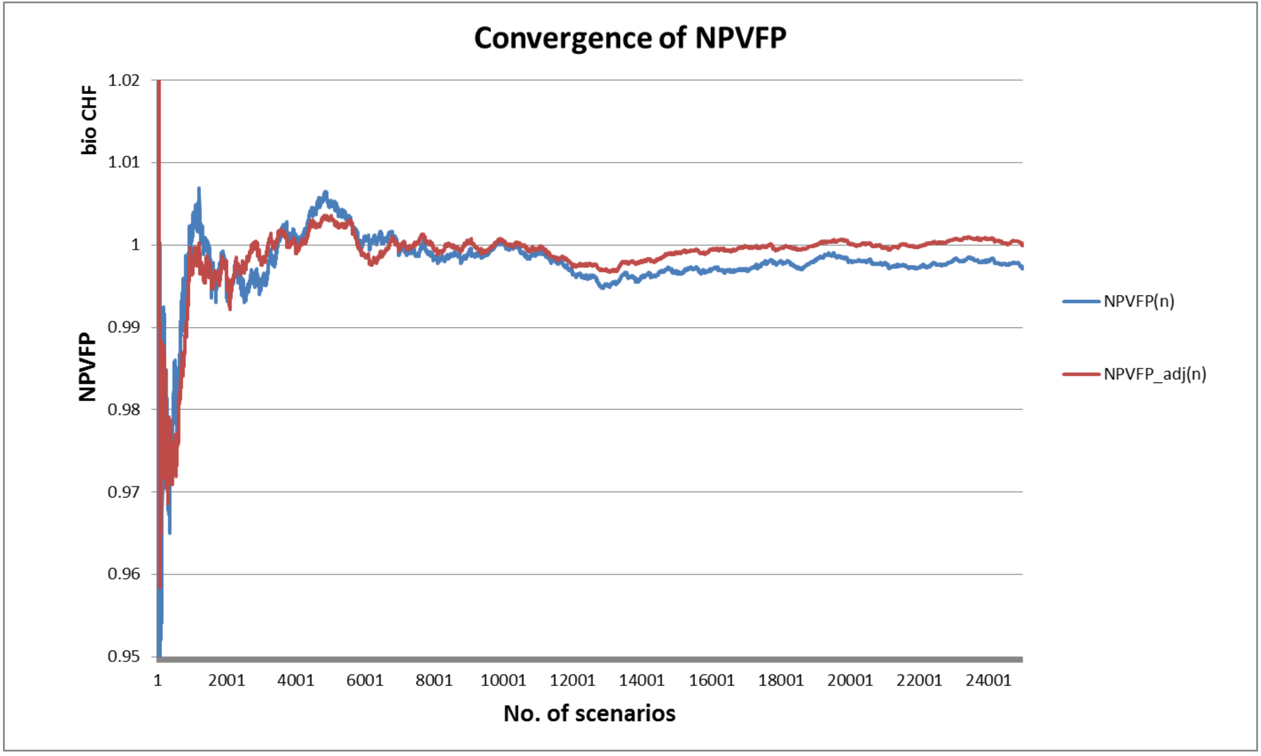

The convergence is much better from the start than in Case 1. On the one hand, this is due to the variance reduction by means of the antithetic pairing. On the other hand, the choice of numeraire was observed to have a significant impact on the convergence. Correspondingly, the scale on the y-axis is much narrower in Figure 4 than in Figure 2. The improved convergence due to the leakage-based control variate can still be seen but is much less pronounced than in Case 1. Note that the values of NPVFP(n) and $NPVFP_{adj}(n)$ for n=25,000 differ by 2.82 mio CHF. This deviation is due to our assumption that the expected value of the leakage is zero, but the mean value of the leakage over our sample of 25,000 antithetic pairs is -2.41 mio CHF.[5] Since this deviation is scaled with b^* =1.168 (see (2)) we get the overall deviation of 2.82 mio CHF. The predicted correction of 2.82 mio CHF to the NPVFP is well within the acceptable range of the statistical error determined by the standard deviation of the mean estimator of 2.81 mio CHF of the NPVFP. There is a bias correction of this kind also in Case 1, but it is less pronounced than in Case 2: Even though the mean leakage of 4.56 mio CHF over the sample exceeds the -2.41 mio CHF of Case 2, the slope of the regression line of 0.116 is much smaller than the 1.168 of Case 2, resulting in a bias correction of -0.53 mio CHF (as opposed to the 2.82 mio CHF of Case 2).

Without variance reduction, the error estimate for the NPVFP on the whole set of 50,000 scenarios is 3.08 mio CHF. As we have seen, the variance reduction due to the antithetic pairing brings it down to 2.81 mio CHF, and the combined variance reduction results in an error estimate of 2.01 mio CHF. Therefore, we conclude that our control variate approach can significantly enhance the variance reduction provided by the antithetic pairs of scenarios.

[5] The standard deviation of the mean estimate for the leakage over the 25,000 antithetic pairs is 1.68 mio CHF, so the sample leakage of 2.41 mio CHF is still in reasonable accord with our assumption that the leakage is zero.

4 REFERENCES AND CONTACT DETAILS

4.1 References

[GL] Glasserman, P. (2001). Monte Carlo Methods in Financial Engineering, Springer.

[KA] Kalberer T. (2006). Market Consistent Valuation of Insurance Liabilities, Der Aktuar, 1, pp. 7-14.

4.2 Contact details

Dr. Andreas Meister, Baloise Switzerland, Basel

Email: *andreas.meister@baloise.ch*

Dr. Mario Hörig, Senior Consultant, Düsseldorf

Email: *mario.hoerig@milliman.com*

Tigran Kalberer, Principal, Zürich

Email: *tigran.kalberer@milliman.com*