QuantAkt

The use of proxy-models for risk measurement of life insurance portfolios

Brownbag Seminar

Lausanne

January 26th, 2016

Guido Grützner

guido.gruetzner@quantakt.com

Agenda

- Introduction
- Replicating Portfolios
- Least Squares Monte Carlo
- Validation
- Applications

Solvency rules, market valuation and proxy-models

- Modern solvency regimes (such as Solvency II or SST) assess risk using market values or market consistent values. This talk will not discuss the details of market consistent valuation and we simply assume:

  The market consistent value at time t (MCV_t) of an insurance portfolio is:
  - The expected value of the discounted cash-flows from the portfolio
  - under a risk neutral (RN) measure calibrated to market conditions at time t

- MCV_0 at the valuation date ($t = 0$) is one of the main components of available capital under solvency frameworks and is nowadays routinely calculated
- Risk is measured by the change of MCV over a time horizon due to risk factors.
  - Risk factors at the end of the time horizon (time $t = 1$) under a real world measure: X
  - Apply a risk measure to the random variable $MCV_1(X)$
- Proxy-Models are currently the standard way to estimate the function MCV_1 for life-insurance portfolios
  - Proxy Models are not concerned with the real world distribution of X
  - Once MCV_1 is “known”, risk measures of $MCV_1(X)$ are calculated from the empirical distribution obtained by simulation

Insurance cash-flows and models

- To determine the MCV of an insurance portfolio (or a whole company) one needs
  - A risk neutral measure, which we will simply assume to be given
  - and the cash flows into and out of the portfolio.
- The cash-flows of life-insurance portfolios are complex, and depend - often in non-linear fashion - on a multitude of underlying risk-factors. Examples:
  - Large number of products (E.g. 50 years of history) with a variety of insurance and savings components
  - Benefits depend on past experience, i.e. are path-dependent
  - Policyholders have a variety of exercise options (such as pay-up or surrender)
  - The investments underlying reserves and the savings processes are a managed multi-asset pool
  - Management has discretion in awarding benefits, i.e. model must contain rules for future management behaviour
  - Projection horizon determined by tenor of the products, i.e. 40 to 60 years
- Cash-flows are not given as explicit functions but require large scale computer models incorporating all of the above features.
  - Building a cash-flow model of an insurance company is a major undertaking (~ 2 to 10 person years)
  - Running it reliably in tight reporting schedules requires highly controlled and optimised processes

Calculation of MCV_0

- The industry standard for calculation of MCV_0 is by Monte-Carlo
  - Simple concept: You only need to feed samples from the RN-measure into your computer model and collect results
  - Convergence is guaranteed and precision is tuned easily by size of samples
- Typical parameters are
  - Sample size between 1000-5000 scenarios with Monte Carlo error around 1%-5%
  - Time required for the run itself: «A few hours»

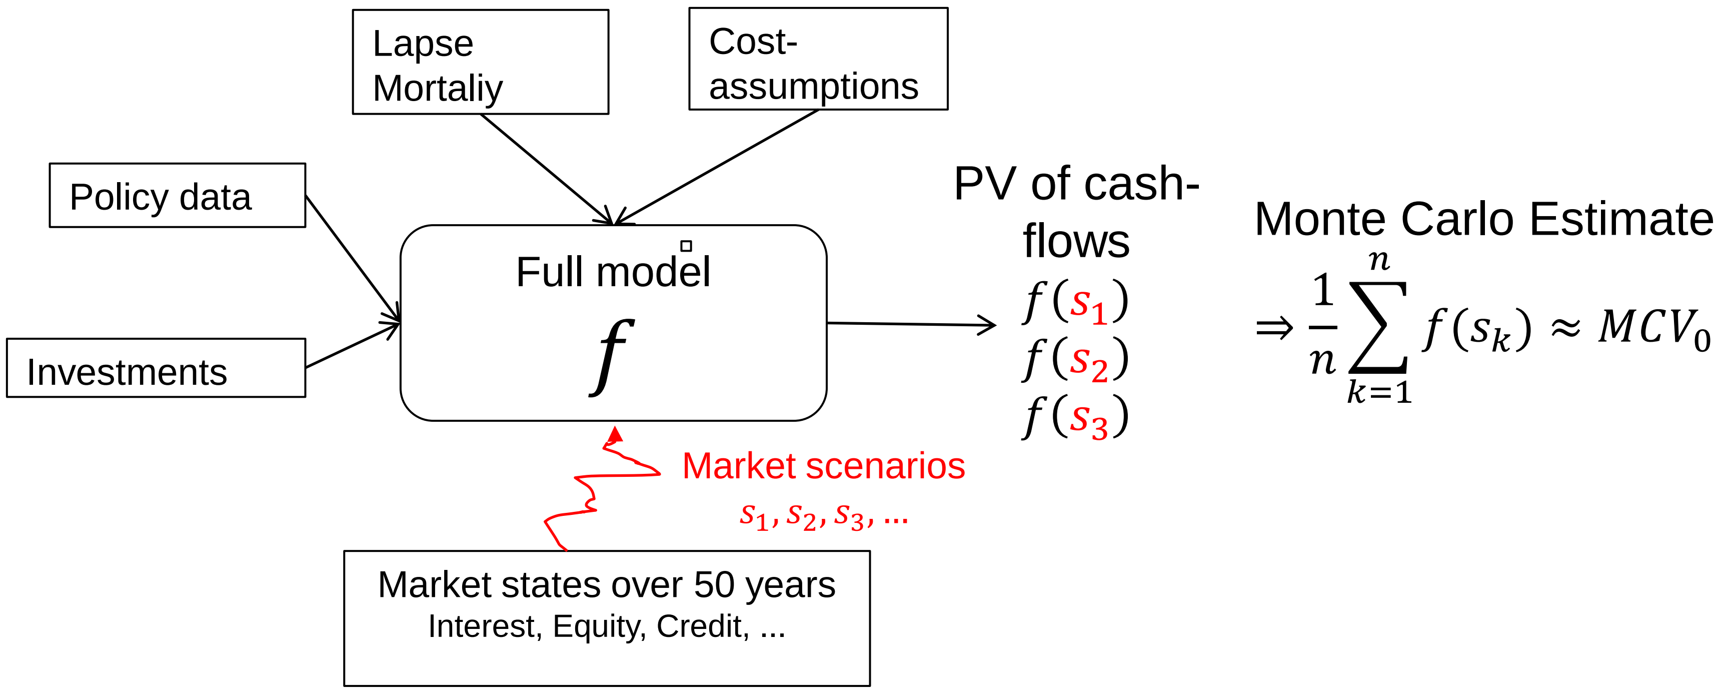

The challenge of "nested stochastics"

- Question: Does Monte Carlo work for MCV_1?
- Answer: In theory yes, BUT ...
  - That would require A LOT OF risk-neutral scenarios
  - and all of those would need to be calibrated to the respective market states at $t = 1$

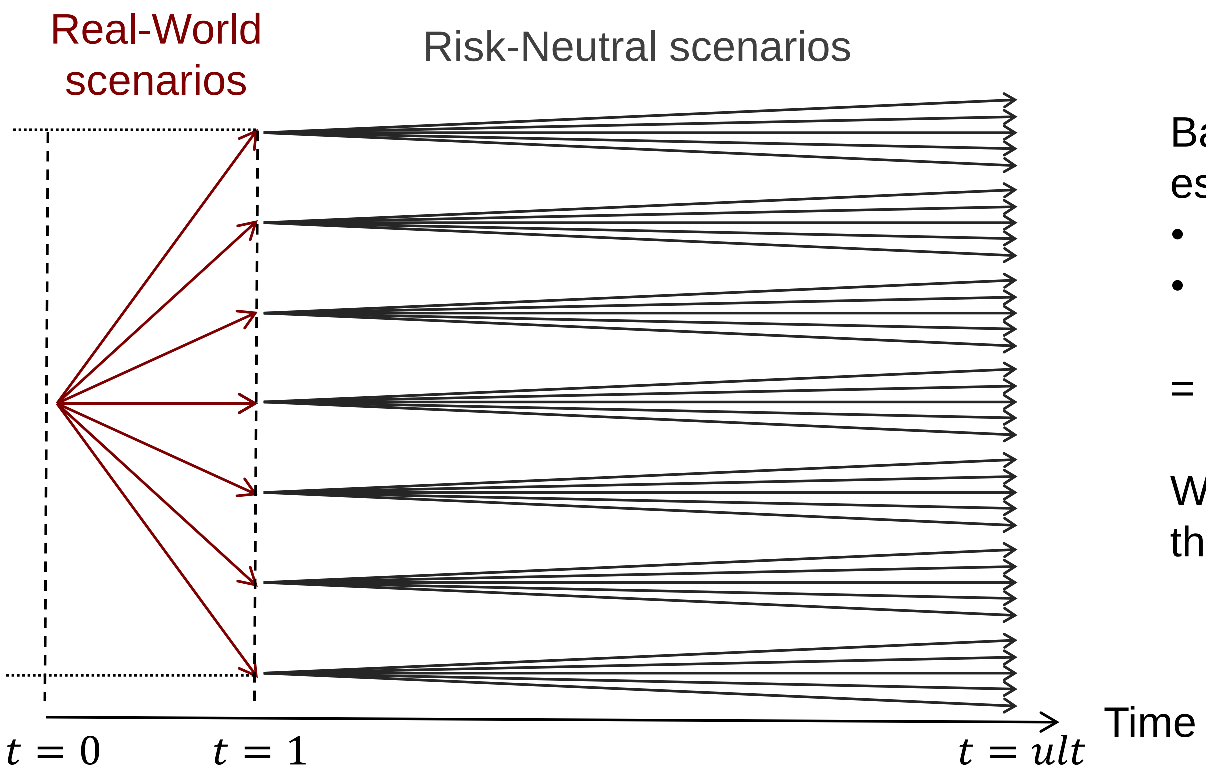

Back-of-the-envelope estimate:

- 10'000 real-world
- each with 1'000 risk-neutral

= 10m projections

With a projection per second this would take 116 days

Enter Proxy-Models

- Proxy-Models are fast approximations to slow and complex computer models
- The idea of a Proxy-Model is "too good to be invented just once"
  - Not restricted to applications in life insurance in other areas known as "surrogate" or "meta" models
  - Applications: Automotive, Aerospace, Geology, Meteorology
- Current status of Proxy-Models in the life insurance industry
  - (Almost) All life insurers with an internal model use a Proxy-Model
  - Currently three approaches : Replicating Portfolios, Least Squares Monte Carlo and Curve Fitting
  - Replicating Portfolios by far the most popular approach (Exception UK: Curve Fitting)
  - Still mainly used for calculation within Solvency II (Solvency Capital Requirement) or SST (Target Capital)
- Still only limited public or academic analysis
  - Mainly presentations by consultants or users with little disclosure of methods and hard facts
- Working papers by the German Actuarial Society (DAV) and the Institute and Faculty of Actuaries
  - links to papers are provided in the Appendix
  - Not really "best practice" papers more "current practice"

- The presenter was member of the DAV working party but takes full responsibility for all statements made in this talk

Agenda

- Introduction
- **Replicating Portfolios**
- Least Squares Monte Carlo
- Validation
- Applications

Replicating Portfolios – a simple recipe

- Three ingredients are required for a Replicating Portfolio

1. A *Design* for the calibration: A set of N scenarios (i.e. paths of the underlying risk factors) $(s_\nu),\ \nu = 1, \ldots, N$
2. A *Universe of Replicating Instruments:* A set of M functions of the scenarios: $(g_\mu),\ \mu = 1, \ldots, M$
3. A *Distance Function:* $d: \mathbb{R}^N \times \mathbb{R}^N \longrightarrow \mathbb{R}^{\geq 0}$
   - Typically Euclidean distance

- The full model f as well as all replicating instruments (g_μ) are evaluated on the scenarios
  - $\left(f(s_\nu)\right)$ as well as each $(g_\mu(s_\nu))$ are N-dimensional vectors
- The replicating portfolio g^* is the linear combination $g^* = \sum_\mu \beta_\mu^* \, g_\mu$ with minimal d –distance to f
- The approximation to the MCV function of f is the MCV function of the replicating portfolio
  - $MCV(f)(x) = E[\, f \,|\, \mathrm{X} = x \,] \approx E[\, g^* \,|\, \mathrm{X} = x \,] = E\left[\, \sum_\mu \beta_\mu^* \, g_\mu \,\middle|\, \mathrm{X} = x \,\right] = \sum_\mu \beta_\mu^* E\left[\, g_\mu \middle|\, \mathrm{X} = x \,\right]$
- To make this work the MCV-function of the replicating instruments must be simple to compute
  - In practice the (g_μ) are simple derivatives such as calls or swaptions with their Black-Scholes value functions

Replication – a toy example

- A single underlying: «annualised return»
- The scenarios are blue and chosen as an iid sample
- The function values are orange and depend in a non-linear way on the underlying

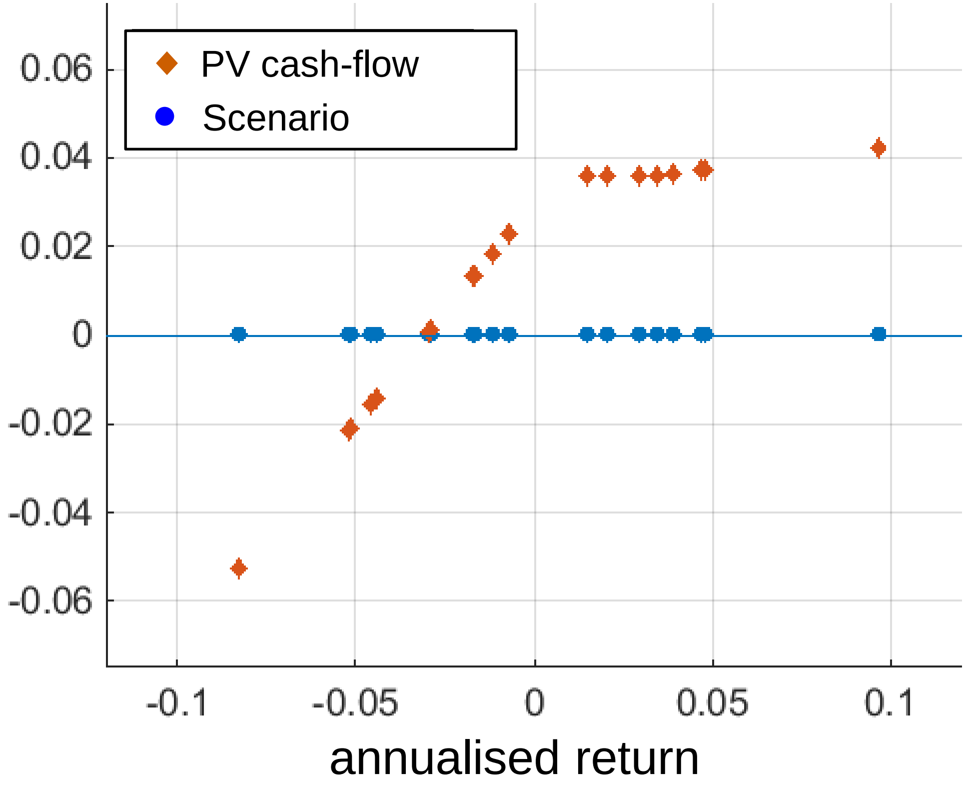

Replication – a toy example

- Replicating instruments are cash, the underlying and calls with various strikes
  - How many instruments? Depends on the shape of the graph
  - Which strikes? Should be in the “interesting” parts of the graph
  - Note that replication is always static not dynamic

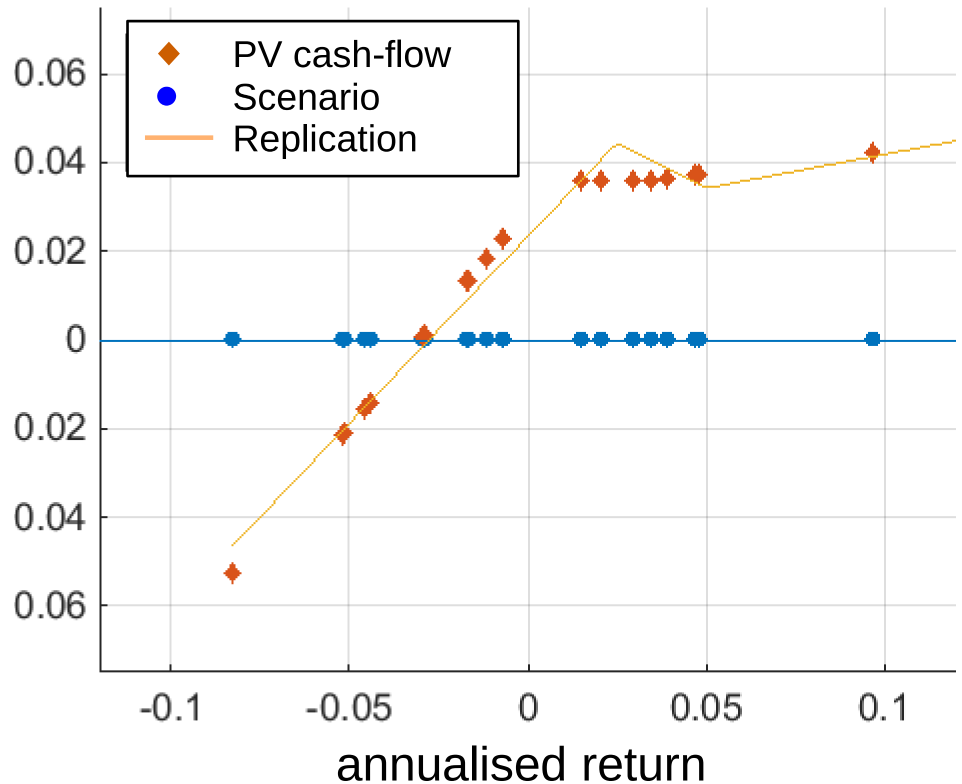

Replication with 4 instruments

- cash,
- the underlying
- call with strike 2.5%
- call with strike 5%

Distance is Euclidean «least squares» fit.

Replicating Portfolio is:
-0.82 cash + 0.85 underlying
- 1.26 call(2.5%) + 0.56 call(5%)

The challenges of replication

- Replicating Portfolios are a recipe indeed: You don’t need any assumptions on the other hand you can’t draw any conclusions!
  - There is obviously no assurance about the approximation quality
- Over time two major challenges have emerged: The question of extrapolation and the choice of replicating instruments
- A variety of techniques have been developed to cope with these challenges
- Both issues are either directly caused by or at least severely aggravated by the “curse of dimension”. The dimension of the approximation problem is in the hundreds and thousands easily.
  - Typical models have between ten and fifty major risk factors or economic variables
  - But the impact of risk factors on cash-flows will depend on the time of impact
  - Each risk factor is a time series of up to 40 or 60 years
  - So the typical dimension is on the order of 10*40 = 400 to 50*60 = 3000
- The problem is not hopeless since not all of those dimensions are equally important!

The first challenge of replication: Extrapolation

- What happens in areas where there are no scenarios for calibration?
  - Can be about “Gaps” in between scenarios or about asymptotic behaviour
  - Extrapolation outside observed values is an implicit result of the calibration, hence difficult to control
  - Since the integration involved in the value calculation is over the whole space even asymptotic behaviour might make a difference
- Practical solution: Avoid extrapolation to the extent possible by carefully choosing the Design for calibration
  - Earliest attempts: Design based on scenarios calibrated to current market conditions (the t=0 scenarios)
  - Next steps and current practice: Include a variety of “stressed” scenarios calibrated to very different market conditions.
  - Typical examples of stresses: Initial yield curve shifts (+/- 100bps) and/or changes in shape
  - The DAV working paper contains a whole section dedicated to this question
- Current state of the art: Designs with total number of scenarios between 5’000 and 50’000 calibrated on up to 50 initial conditions
- Future approach? Space filling Designs – to the extent possible
  - This is obviously difficult given the curse of dimension
  - More than 50 (say) initial calibrations pose similar technical difficulties as nested stochastics
- Current Designs are a compromise between mathematical and practical requirements

The second challenge: Choice of instruments

- There are two very different approaches to the choice of instruments
  - Manual selection based on experience and prior knowledge
  - Automatic selection based on a selection algorithm
- The two methods have clear and distinct pros and cons
- Manual selection:
  - Time consuming, especially after cash-flows changed a lot (due to changes in markets, the portfolio or the model)
  - Manual RPs will have a limited complexity, i.e. the number of instruments is limited to a few tenths or (low) hundreds
  - The approach is subjective and the final quality may vary
- Automatic selection
  - For this to work you need a large number of instruments. Due to the curse of dimension, “large” becomes quickly VERY LARGE
  - Naïve approach quickly leads to problems, e.g. more instruments than scenarios, high correlation and close-to-null-portfolios with ensuing numerical instability,
  - Hence more advanced statistical methods are required such as feature selection, regularisation techniques
- Mixtures of the two approaches are possible and used in practice

Agenda

- Introduction
- Replicating Portfolios
- **Least Squares Monte Carlo**
- Validation
- Applications

Least Squares Monte Carlo (LSMC)

- Least Squares Monte Carlo addresses the big weakness of nested stochastics:The large number of risk-neutral scenarios

Nested Stochastics

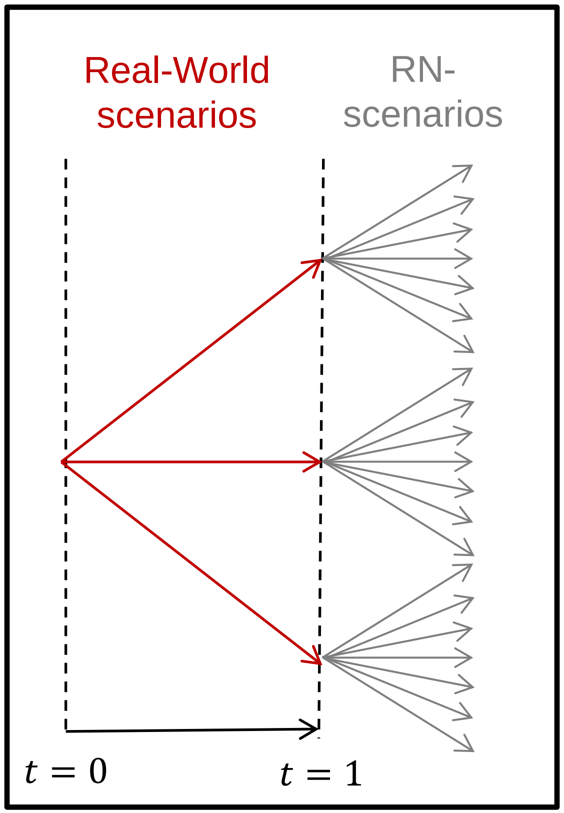

Precise estimate due to MANY risk-neutral scenarios

LSMC

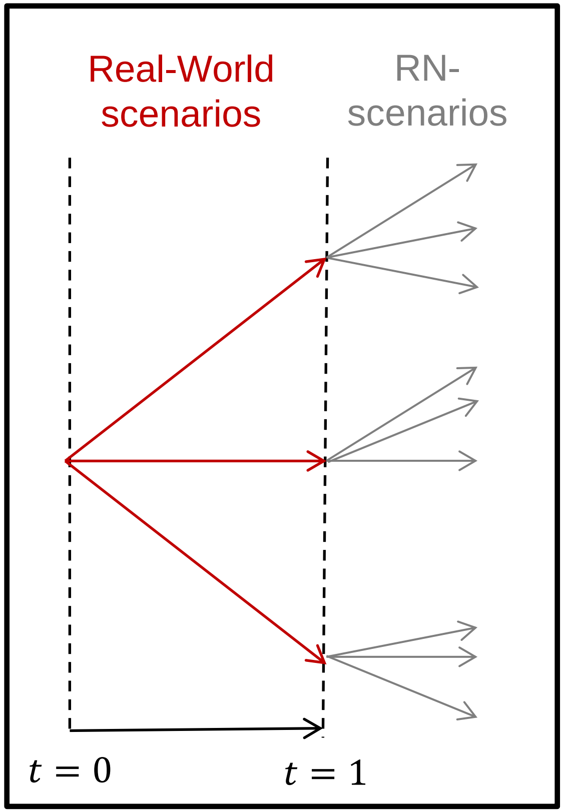

Precise estimate due to averaging out of the Monte-Carlo Error

How does “averaging-out the error” work?

- The error of a Monte-Carlo estimate with few risk-neutral scenarios fulfills (quite precisely) the assumptions of regression:
  - The error for each Real-World scenario has (asymptotically) a normal distribution
  - The errors are independent for different Real-World scenarios

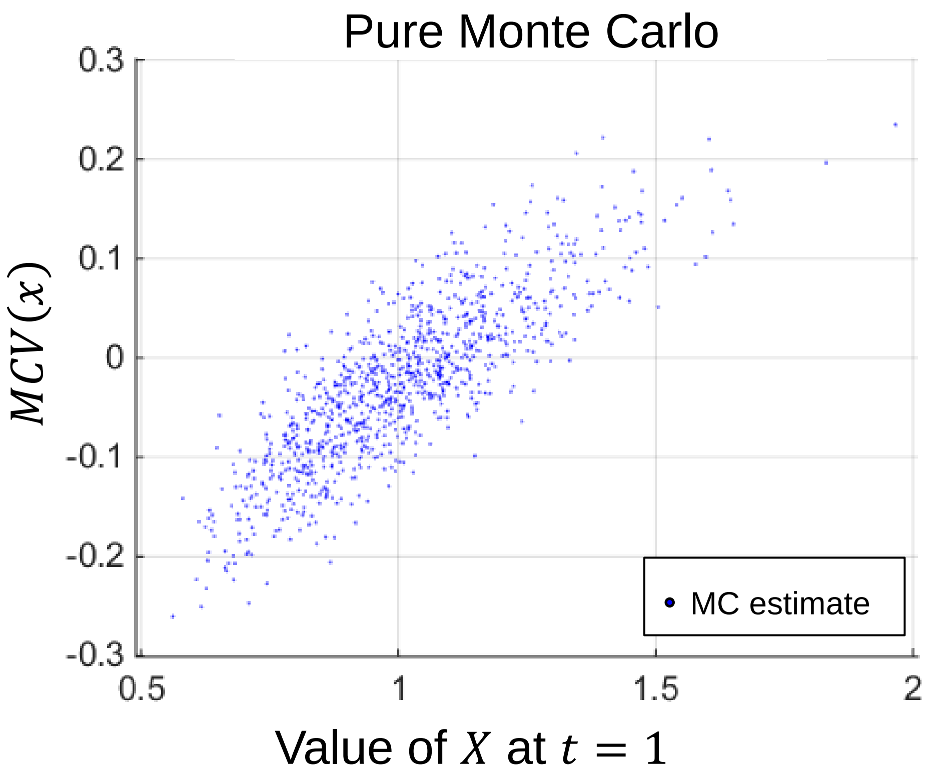

- MC-Estimates calculated here using 1000 real-world each with 100 risk-neutral scenarios
- The low number of risk-neutral scenarios produces large MC-errors and broadly scattered estimates

How does “averaging-out the error” work?

- A “Least Squares Monte Carlo” Proxy-Model is just multivariate regression
  - This makes all tools from regression available
  - The remaining challenge is the generation of nested scenarios

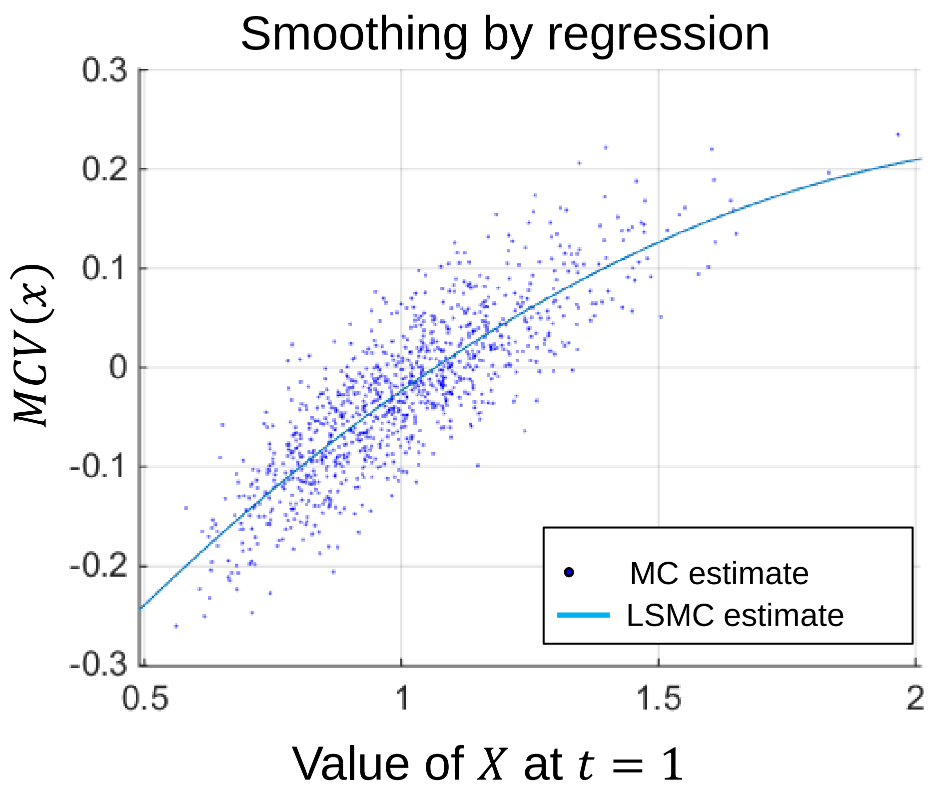

- In the example the regression function is a quadratic polynomial
- This regression function is the estimate of the MCV-function at $t = 1$.

Comparison of the two methods

- The table compares some crucial aspects of the two methods

	Replicating Portfolio	**Least Squares MC**
Risk factors	Mainly market risks	All modelled
Stat. assumptions	None	Regression assumptions
Basis functions	Financial instruments	General functions/polynomials
Cross effects	In principle	possible
Use	Widespread for many years	Growing interest
Dimension	Very high (hundreds, thousands)	High (several dozen)
Design-Scenarios	Close to MCEV-Scenarios 10 to 50 RW x 500 to1000 RN	Special nested scenarios 5000 bis 20’000 RW x 2 RN

Agenda

- Introduction
- Replicating Portfolios
- Least Squares Monte Carlo
- Validation
- Applications

Validation of Proxy-Models

- Neither Replicating Portfolios nor LSMC models come with a guarantee of approximation quality. This makes extensive validation necessary.
- The validation methods are quite similar for both approaches:
  - In sample and out-of sample measures of fit (R^2, quadratic error)
  - Targeted analysis of residuals: Are there certain risk factors or certain areas where fit is especially bad? Does the shape of residuals suggest improvements?
- Analysis of the sensitivity of results to inputs or parameters of the methods used
  - Choice of financial instruments or regression functions
  - Change of the Design for calibration
  - Stability over time
- Validation includes qualitative aspects as well
  - Is the MCV function or the Replicating Portfolio plausible?
  - Is it consistent with a-priori knowledge?

Validation by calculation

- The most convincing test is the comparison of the Proxy-Model with the full model for certain sensitivities, i.e. market states at $t = 1$
- Typical test cases are values in the relevant tail of the MCV distribution
- Test can be done for single and combined stress-scenarios
  - But curse of dimension applies again
- One method can be used to validate the results of the other method
- The number of test cases is limited by the overall computational budget

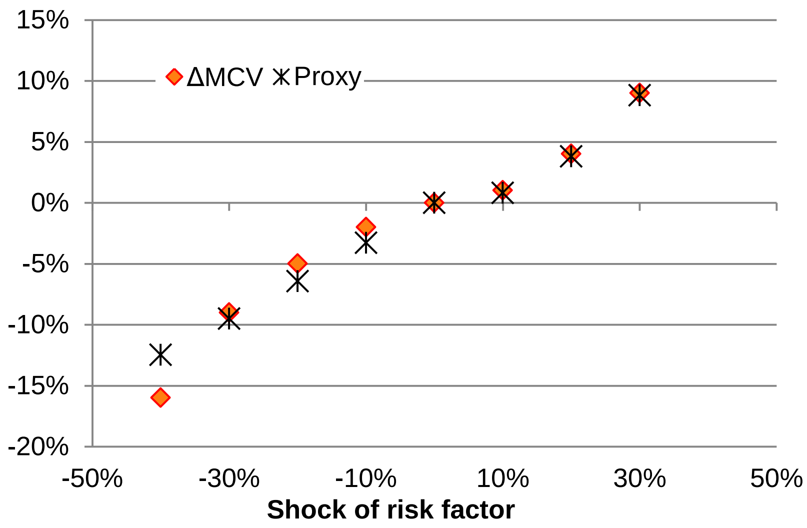

- The example shows a stress for a single risk factor
- Delta MCV is calculated with the full model
- The example shows strong deviations for downward stresses of the risk factor

Agenda

- Introduction
- Replicating Portfolios
- Least Squares Monte Carlo
- Validation
- Applications

Applications: Analysis of the full model

- A Proxy-Model is arguably the only way to understand the full model
  - Only by calculating and comparing multiple risk factor combinations and parameters is it possible to get a clear picture what is going on

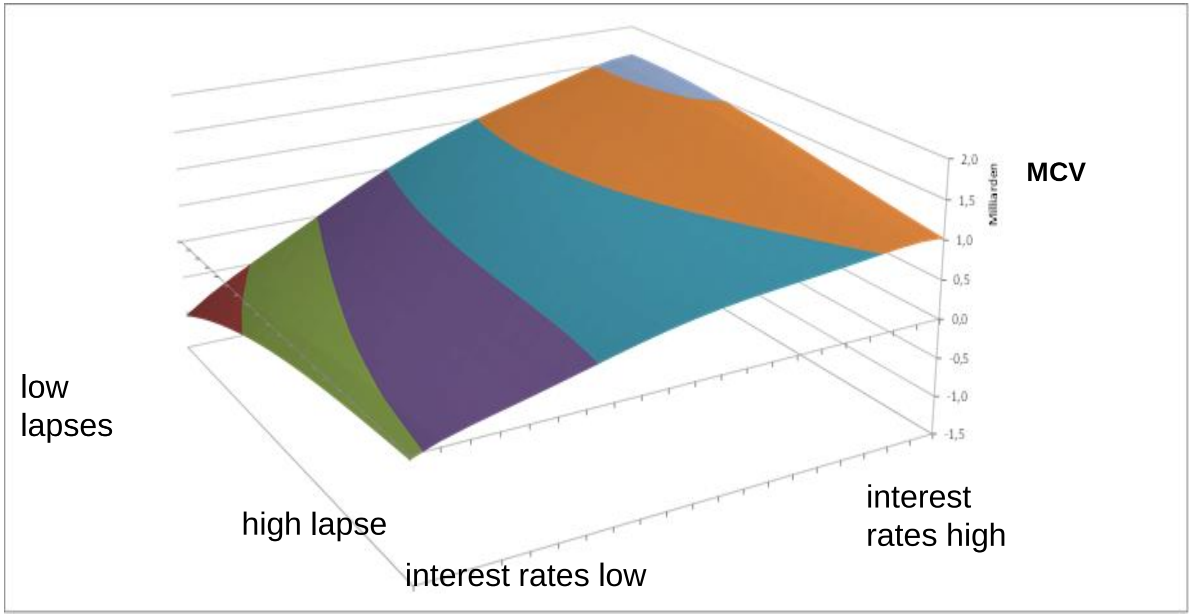

The graph shows the combined effect of lapse assumptions and long-term interest rates

Specific questions about the full model

- "Understanding the full model" is not an theoretical or academic exercise !
- Very practical questions can be asked and the answers have practical consequences
- Examples:
  - What risk factors are "important", i.e. have strong impact, which are less important?
  - Which risk factors are linear and isolated from other risk-factors?
  - Which risk-factors have non-linear impact or show strong interactions?
- Practical consequences:
  - "Important" and "unimportant" is obvious, isn't it?
  - Linear and isolated risk factors are good candidates for cheap and simple hedging
  - Non-linear and strongly interacting risk factors are difficult to hedge

Typical current applications of Proxy-Models

- Currently the main application is the calculation of solvency capital under SST or Solvency II over the one year horizon at a valuation date
- Additional applications are
  - Projections and forecasts, fast close
  - Quarterly solvency assessments
  - Capital budgeting and risk appetite assessments
- These additional applications do not require new methods but operationally robust processes
  - The Proxy-Model needs to be calibrated and validated several times during the year

Advanced applications

- Advanced applications of Proxy-Models require methodological improvements
- “As-if” calculations which depend on parameters
  - Example: SST requirement as a function of the amount of Euro-investment of a Swiss life insurer
- Questions which require optimisation
  - What is the “optimal” amount of Euro-investment?
  - Optimisation is difficult since one is not calculating a single MCV function but looking for “the best”
- Questions about time development and dynamic strategies
  - The market consistent value is not just a single number MCV_1 but a process in time
  - You might want to understand and control MCV_2, MCV_3, … as well
- The most prominent field for advanced applications is in Asset Liability Management
  - Strategic Asset Allocation: Optimal high-yield and foreign investments
  - Impact of hedging strategies on solvency requirements
  - Optimal dynamic hedging strategies incorporating aspects like basis risk, liquidity

Proxy-Models in a wider context

- While the origin and main purpose of Proxy-Models is calculation of solvency capital the methods and problems apply to a wider context
  - The basic problem is understanding a non-linear and high-dimensional function from observed samples
- Such and similar questions are applicable to many areas
  - Prediction of client behaviour based on past experience and demographic data
  - Automated underwriting based on prediction of claims frequency and severity
- This problem and the methods for its solution are the typical topics of machine learning
  - Regression, factor selection, regularisation,……
- A thorough understanding of these methods has wide applicability beyond Proxy-Models and is an opportunity for actuaries to apply their knowledge to new and rewarding problems

Appendix

Links to the working papers

- The working paper of the German Actuarial Society (in German!):
  - https://aktuar.de/unsere-themen/fachgrundsaetze-oeffentlich/2015-07-08_DAV_Ergebnisbericht%20AG%20Aggregation.pdf
- The working paper of the Faculty and Institute of Actuaries
  - http://www.actuaries.org.uk/research-and-resources/documents/heavy-models-light-models-and-proxy-models-working-paper