



Monte Carlo Methods for Non-Life Actuaries

Reserve Risk, Aggregate Losses, and Capital Allocation in R

Sigrid Norrgård



ODIN PRESS



Monte Carlo Methods for Non-Life Actuaries

Copyright © 2026 by **Sigrid Norrgård**.

All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopying, recording, or otherwise—without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law.

Paperback version, First edition: 2026.

ISBN: forthcoming (Bowker allocation pending)

Trim Size: 6.14 × 9.21 inch.

Source-code copyright. All R source code listings reproduced in this volume—including those in the Appendix, the chapter bodies, and any auxiliary file distributed alongside this book—are the copyrighted intellectual property of **Sigrid Norrgård**. No reproduction, redistribution, adaptation, translation into another programming language, or incorporation into any derivative work, whether academic, commercial, or otherwise, is permitted without the prior written permission of the copyright holder. Single-machine execution by the individual purchaser of this volume for the sole purpose of reproducing a figure or table from this book is granted as a limited, non-transferable, non-sublicensable use and does not constitute redistribution.

Published by Odin Press.

<https://odinpress.org>

This book was typeset in L^AT_EX by Odin Press from a manuscript submitted by the author.

Printed and Bound in the UNITED STATES OF AMERICA.

9 8 7 6 5 4 3 2 1 0

Preface

Non-life actuaries have read Glasserman’s “Monte Carlo Methods in Financial Engineering” for two decades and learned variance reduction, low-discrepancy sequences, and the MCSE-based reporting discipline that the book established as the field’s professional standard. They have read Mack’s 1993 ASTIN paper through the Wüthrich–Merz monograph and absorbed the analytic prediction error for the chain ladder. They have read Klugman, Panjer, and Willmot’s “Loss Models” and absorbed the compound-distribution machinery. The three traditions sit on a non-life reserving actuary’s desk, and each is excellent on its own ground.

Each is also incomplete for the question a non-life internal-model team is actually asked at year-end. Glasserman is calibrated against derivatives books, not motor TPL triangles. Mack is the analytic core but is silent on the catastrophe overlay, the copula coupling, the reinsurance treaty layer, and the capital allocation. Klugman covers the compound aggregate but stops short of the Solvency II one-year horizon and the regulator-facing audit defense. The actuary who wants to ship a defensible internal-model submission must currently stitch the three traditions together from scratch, inventing the methodology disclosure and the diagnostic suite anew at every annual recalibration.

This book is that stitch, written explicitly for the non-life setting. It covers the full pipeline from per-claim severity to the firm-level economic capital model, with the Solvency II one-year horizon as the binding regulatory framework throughout. Every numerical result in the chapters is produced by R code printed in Appendix A and reproducible by the reader; every methodology choice is named, defended, and routed against the regulator’s likely audit response. The audience is the practitioner who already knows variance reduction and MCSE — and who wants the non-life specialisation the existing literature does not provide.

The book delivers five things the existing literature does not assemble together in a single defensible pipeline. First, the reserve-risk machinery on

real triangles: bootstrap chain ladder, Mack analytic, GLM Tweedie, and Bayesian hierarchical, with the cross-method reconciliation that makes the twenty-to-forty-percent gap between bootstrap SD and Mack analytic prediction error visible rather than buried. Second, the frequency–severity–tail engine: Poisson and negative-binomial count models, Panjer recursion, peaks-over-threshold with the GPD threshold-selection discipline a regulator will audit, and the spliced and generalized-hyperbolic alternatives when POT fails. Third, the multi-line aggregation: t-copula and R-vine families calibrated against the joint-extreme evidence the calibration window allows, with the EVT-heuristic bound that quantifies what the data cannot anchor. Fourth, the reinsurance and catastrophe machinery: layer pricing under the Mildenhall decomposition with Bühlmann–Straub credibility on thin-experience layers, and the ELT-driven cat module with the secondary-uncertainty discipline that catches the demand-surge double count my Munich Re years taught me to refuse. Fifth, the Solvency II finishing pipeline: Merz–Wüthrich one-year reserve risk, two-state Markov underwriting cycle for premium risk, Euler and Aumann–Shapley capital allocation under coherent risk measures, and the full economic capital model on a four-line plus cat-overlay portfolio.

The chapters proceed in five movements. Chapter 1 fixes the Monte Carlo setting — LLN, CLT, MCSE, RNG discipline — that every subsequent chapter relies on. Chapters 2 through 5 are the reserving toolkit: bootstrap chain ladder, Mack’s distribution-free model, GLM stochastic reserving across the Tweedie family, and the Bayesian MCMC interlude with Hamiltonian Monte Carlo and the no-U-turn sampler. Chapters 6 and 7 are the frequency–severity engine: Poisson and negative-binomial count models with Panjer aggregation; peaks-over-threshold severity tails with spliced and generalized-hyperbolic alternatives. Chapter 8 calibrates the t-copula. Chapter 9 prices reinsurance layers. Chapter 10 integrates the vendor catastrophe model. Chapter 11 produces the aggregate portfolio loss distribution with a three-method cross-validation (simulation, FFT, Panjer) that audits the tail. Chapters 12 and 13 deliver the Solvency II reserve-risk and premium-risk modules. Chapter 14 develops capital allocation. Chapter 15 assembles everything into the firm-level economic capital model on the calibrated multiline_casualty and property_continental cat-overlay portfolio.

Appendix A prints the decisive R script for each chapter — fifteen scripts that together reproduce every printed numerical result. Appendix B carries one hundred and five applied exercises in seven exercises per chapter, organised by chapter and covering the methodology choices a non-life actuary faces in the year-end review and the regulatory submission. Both appendices are

integral to the book; the chapters develop the machinery, the appendices show the reader how the machinery is run and how it is defended.

The book's defensibility-first orientation is the orientation I have brought to twenty-six years of year-end sign-offs across the Nordic primary and reinsurance markets — eleven at If P&C, twelve at Munich Re Stockholm, and the last three as an independent consultant on Solvency II internal-model recalibrations. Every methodology choice in the chapters is named, every diagnostic is run, and every result that lands in the SCR submission is reproducible from the code in Appendix A. The book does not pretend that Monte Carlo simulation produces capital numbers without methodology commitments; it makes the commitments explicit and shows the reader how to defend them when the Finansinspektionen, BaFin, or the PRA asks.

Sigrid Norrgård. Saltsjöbaden, 2026.

Contents

Preface	vii
List of Figures	xvii
List of Tables	xxi
1 The Actuarial Monte Carlo Setting	1
1.1 The Monte Carlo estimator, its SLLN, and the MCSE obligation	2
1.2 L’Ecuyer-CMRG and substream-safe Monte Carlo on the motor_nordic baseline	7
1.3 Antithetic variates: theory, R implementation, and measured variance reduction	12
1.4 Control variates and stratified sampling as orthogonal levers	16
1.5 CRN-paired comparison across the three variance-reduction techniques	20
2 Bootstrap Chain Ladder and ODP	27
2.1 Chain ladder development factors and the ODP exponential family	28
2.2 ODP bootstrap consistency and the BCL-ODP moment equivalence	32
2.3 Bootstrap algorithms — BCL, ODP, capped residuals, and tail-factor extension	37
2.4 Reserve distributions on motor_nordic (long-tail) and property_continental (short-tail)	43
2.5 Residual diagnostics and the process-vs-estimation error decomposition	49
3 Mack’s Distribution-Free Model	55

3.1	Mack’s first- and second-moment assumptions and the age-to-age sigma estimator	56
3.2	Mack’s closed-form MSE formula and its second-moment recursion proof	60
3.3	Parametric bootstrap for Mack VaR/TVaR	63
3.4	Multi-line Mack with Quarg-Mack extra-trend and sigma-floor on thin tails	68
3.5	Mack vs BCL bootstrap — analytic-simulation equivalence on the same portfolios	73
4	GLM Stochastic Reserving — ODP, Gamma, Tweedie	79
4.1	Exponential family, dispersion, and the GLM link function	80
4.2	Tweedie compound Poisson-Gamma and the p-parameter boundary	83
4.3	GLM reserve consistency and the sandwich variance under misspecification	86
4.4	Algorithms — IRLS, profile likelihood for Tweedie p, and sandwich bootstrap	89
4.5	Reserve distributions on property_continental and motor_nordic across families	93
5	Bayesian Reserve Models — the MCMC Interlude	101
5.1	Hierarchical Bayesian reserve model and the actuarial-notation posterior	102
5.2	Metropolis-Hastings — acceptance ratio, ergodicity, and reserve-parameter sampling	105
5.3	Hamiltonian Monte Carlo — leapfrog integrator and NUTS adaptation	107
5.4	Convergence diagnostics — ESS, R-hat, and the posterior predictive check	110
5.5	Posterior reserve distributions on motor_nordic and multi-line_casualty	113
6	Claim Frequency Models — Poisson, NB, and Panjer	123
6.1	Poisson process, NB-as-Gamma-mixture, and the (a,b,0) class	124
6.2	Panjer recursion theorem and its generating-function proof	129
6.3	MLE for Poisson and NB count data with Cameron-Trivedi over-dispersion	131

6.4	Algorithms — Poisson thinning, NB Gamma-Poisson decomposition, Panjer compound	135
6.5	property_continental over-dispersion against the calibrated $\text{phi_freq} = 2.5$	137
7	Severity Tail Modeling — POT, GPD, Spliced, GH	145
7.1	Peaks-over-threshold and the generalized Pareto distribution	146
7.2	Hill estimator consistency and GPD MLE asymptotic normality	149
7.3	Threshold selection via mean-excess plot, Hill plot, and stability diagnostics	152
7.4	Spliced lognormal-GPD and generalized hyperbolic body-and-tail families	155
7.5	multiline_casualty PI+BI severity tail vs the calibrated $\text{xi}=0.28$	159
8	Copula-Coupled Multi-Line Aggregation	167
8.1	Copula definition, Sklar’s theorem, and the Gaussian-vs-t-copula contrast	168
8.2	Tail-dependence: Gaussian zero, t-copula $\text{lambda_U}(\text{rho},\text{nu})$, and the Archimedean family	172
8.3	Aggregate VaR inflation under t vs Gaussian dependence .	175
8.4	Algorithms — Cholesky, chi-squared-mixture, Archimedean generator, and R-vine construction	178
8.5	Four-line aggregation on multiline_casualty under Gaussian, t, Clayton, and R-vine	182
9	Reinsurance Layer Pricing	191
9.1	Layer mechanics, pure premium, risk loading, and reinstatement premium	192
9.2	Layer premium decomposition (Mildenhall) and Buhlmann-Straub credibility	195
9.3	Algorithms — per-occurrence XL, aggregate XL with reinstatements, stop-loss, surplus	197
9.4	Stochastic XL pricing on multiline_casualty PI+BI severity tail	200
9.5	Credibility, surplus shares, and per-line vs aggregate stop-loss cost	204
10	Catastrophe Modeling and Vendor Integration	211

10.1	Event-loss-table, AEP, OEP, and the secondary-uncertainty CV	212
10.2	ELT-integrated AAL convergence and Bayesian model averaging proof	215
10.3	Algorithms — Poisson event simulation, secondary uncertainty, AEP/OEP, vendor blending	217
10.4	property_continental windstorm AAL via ELT simulation	219
10.5	Two-vendor blending and secondary-uncertainty sensitivity sweep	224
11	The Aggregate Portfolio Loss Distribution	231
11.1	Collective risk model, $S = \sum X_i$, and the characteristic function machinery	232
11.2	FFT inversion for compound Poisson density and existence under finite mean	235
11.3	TVaR coherence and the Acerbi-Tasche proof	237
11.4	Algorithms — direct simulation, FFT inversion, and Panjer aggregation	239
11.5	Three-method cross-validation on multiline_casualty + cat overlay	242
12	One-Year Reserve Risk under Solvency II	255
12.1	One-year vs ultimate reserve risk, the CDR, and SCR_reserve	256
12.2	Merz-Wuthrich one-year volatility formula and its proof .	259
12.3	MW vs parametric-bootstrap equivalence in the large-triangle limit	263
12.4	Algorithms: MW analytic, CDR-bootstrap, per-line MW, and diversification matrix	265
12.5	motor_nordic and multiline_casualty one-year SCR_reserve	270
13	Premium Risk and the Underwriting Cycle	279
13.1	Premium risk, underwriting cycle, and the loss-ratio decomposition	280
13.2	UW-cycle stationary distribution via Perron-Frobenius . .	283
13.3	Premium risk under cycle = LR variance + state-mixing variance	286
13.4	Algorithms: Markov-cycle simulation, fixed-regime and switching premium risk, MLE	288

13.5	property_continental cycle MLE and premium VaR by cycle state	292
14	Capital Allocation — Euler, Aumann-Shapley, Coherence	301
14.1	Coherent risk measures and the capital-allocation problem	302
14.2	Euler/gradient allocation and the conditional-expectation equivalence	304
14.3	Aumann-Shapley as the unique coherent allocation; TVaR coherence	307
14.4	Algorithms: Monte Carlo Euler-VaR/TVaR, AS line-integral, baselines, MCSE	311
14.5	multiline_casualty 4-line allocation under Euler, AS, and proportional methods	314
15	The Portfolio Capital Model — Assembling the Internal Model	323
15.1	Economic capital model, risk-adjusted capital, and capital efficiency by line	324
15.2	Coherence of the assembled ECM and t-copula sub-additivity	327
15.3	Algorithms: full pipeline, stress shocks, reinsurance optimization, annual update	329
15.4	End-to-end run on multiline_casualty + property_continental cat overlay	333
15.5	Stress scenarios, reinsurance optimization, and capital efficiency ranking	337
A	Source Code	345
A.1	Chapter 1 — The Actuarial Monte Carlo Setting	345
A.2	Chapter 2 — Bootstrap Chain Ladder and ODP	351
A.3	Chapter 3 — Mack’s Distribution-Free Model	360
A.4	Chapter 4 — GLM Stochastic Reserving — ODP, Gamma, Tweedie	368
A.5	Chapter 5 — Bayesian Reserve Models — the MCMC Interlude	375
A.6	Chapter 6 — Claim Frequency Models — Poisson, NB, and Panjer	381
A.7	Chapter 7 — Severity Tail Modeling — POT, GPD, Spliced, GH	383
A.8	Chapter 8 — Copula-Coupled Multi-Line Aggregation	387
A.9	Chapter 9 — Reinsurance Layer Pricing	391

A.10	Chapter 10 — Catastrophe Modeling and Vendor Integration	397
A.11	Chapter 11 — The Aggregate Portfolio Loss Distribution .	402
A.12	Chapter 12 — One-Year Reserve Risk under Solvency II .	407
A.13	Chapter 13 — Premium Risk and the Underwriting Cycle	417
A.14	Chapter 14 — Capital Allocation — Euler, Aumann-Shapley, Coherence	422
A.15	Chapter 15 — The Portfolio Capital Model — Assembling the Internal Model	427
B	Exercises	435
B.1	The Actuarial Monte Carlo Setting	435
B.2	Bootstrap Chain Ladder and ODP	440
B.3	Mack’s Distribution-Free Model	446
B.4	GLM Stochastic Reserving — ODP, Gamma, Tweedie . .	452
B.5	Bayesian Reserve Models — the MCMC Interlude	458
B.6	Claim Frequency Models — Poisson, NB, and Panjer . . .	464
B.7	Severity Tail Modeling — POT, GPD, Spliced, GH	471
B.8	Copula-Coupled Multi-Line Aggregation	477
B.9	Reinsurance Layer Pricing	483
B.10	Catastrophe Modeling and Vendor Integration	490
B.11	The Aggregate Portfolio Loss Distribution	496
B.12	One-Year Reserve Risk under Solvency II	503
B.13	Premium Risk and the Underwriting Cycle	509
B.14	Capital Allocation — Euler, Aumann-Shapley, Coherence .	516
B.15	The Portfolio Capital Model — Assembling the Internal Model	522
	References	529
	Index	537
	About the Author	541

Chapter 1

The Actuarial Monte Carlo Setting

Every reserve \widehat{R} , every capital requirement, every reinsurance price this book reports is the output of a Monte Carlo run, and a Monte Carlo run is only as honest as the discipline the actuary brought to it. Actuarial Monte Carlo is not generic statistical simulation — it is a regulator-facing inferential practice with binding standards on the standard error reported, the random stream that produced it, and the audit trail back to the code that ran. The opening chapter fixes those standards.

The deliverable is a single calibrated baseline for the `motor_nordic` reserve estimator on which the rest of the book reports its measurements. That baseline is the naive Monte Carlo estimator $\widehat{\mu}_N$ run under a parallel-safe L'Ecuyer-CMRG generator with documented substreams; the operational signature is a measured Monte Carlo standard error on every quoted figure, with the headline closing artifact being a variance-reduction factor table that benchmarks antithetic, control-variate, stratified, and common-random-numbers reductions against the unreduced baseline. A reserve quoted without its MCSE envelope is, by my reading after twenty-six years of sign-offs, professional malpractice — and the protocol this chapter sets is what makes the refusal mechanically possible.

The chapter is built in five sections. §1.1 develops the naive Monte Carlo estimator, the strong law and central limit theorem that govern it, and the MCSE reporting protocol every later chapter inherits. §1.2 fixes the L'Ecuyer-CMRG pseudo-random stream and the substream-allocation pattern for parallel work. §1.3 introduces antithetic variates with the monotone-

integrand variance bound. §1.4 develops the control-variate and stratified-sampling orthogonal levers. §1.5 closes with the common-random-numbers paired-comparison driver and the variance-reduction factor table the rest of the book cites.

1.1 The Monte Carlo estimator, its SLLN, and the MCSE obligation

The operational discipline of regulator-grade actuarial Monte Carlo rests on five commitments made once and inherited by every later chapter. Every reported number carries the Monte Carlo standard error that produced it; the pseudo-random stream is a parallel-safe L'Ecuyer-CMRG generator with named substreams, never an unqualified seed on a Mersenne Twister default; variance-reduction levers are deployed on a single named baseline simulator with realized variance ratios *measured* against that baseline rather than asserted theoretically; every method-versus-method comparison runs under common random numbers unless the methods explicitly cannot share inputs; every quoted number traces to a named script in `book_code/` with a documented seed. The baseline portfolio is `motor_nordic` — a long-tail Nordic motor third-party-liability book whose calibration constants live in `book_meta.yaml`. Subsequent chapters of this book refine the same baseline under successively more elaborate reserve estimators; the variance-reduction factor table that closes this chapter is the scaffold against which those chapters report MCSE deltas.

Fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and a real-valued, integrable random variable $X: \Omega \rightarrow \mathbb{R}$ with distribution F . Let $g: \mathbb{R} \rightarrow \mathbb{R}$ be a measurable function with $\mathbb{E}[g(X)] < \infty$. The quantity of interest is

$$\mu = \mathbb{E}[g(X)] = \int_{\mathbb{R}} g(x) dF(x). \quad (1.1)$$

In the `motor_nordic` setting, $X = (Z_1, \dots, Z_{10})$ is the vector of ten log-increments of the cumulative paid loss across development periods, F is the joint lognormal distribution with the calibrated (μ_0, σ) , and $g(X) = \exp(\sum_{j=1}^{10} Z_j)$ is the ultimate loss per accident year. The integral in (??) has no closed form for the joint distribution actuaries actually fit, and so a sample mean is what gets reported.

The naive Monte Carlo estimator of μ , based on N independent and identically distributed replicates X_1, \dots, X_N drawn from F , is

$$\hat{\mu}_N = \frac{1}{N} \sum_{b=1}^N g(X_b). \quad (1.2)$$

The estimator is unbiased by linearity of expectation: $\mathbb{E}[\hat{\mu}_N] = \mu$ for every finite N , provided μ exists. Two non-trivial results govern the rest of this book: the strong law of large numbers, which says $\hat{\mu}_N$ converges to μ almost surely; and the central limit theorem, which says the residual error $\hat{\mu}_N - \mu$ is approximately Gaussian with standard deviation σ_g/\sqrt{N} for all sufficiently large N , where $\sigma_g^2 = \text{Var } g(X)$.

1.1.1 Strong law of large numbers

Theorem 1.1 (SLLN for the Monte Carlo estimator). *Let X_1, X_2, \dots be i.i.d. with distribution F , and let g be measurable with $\mathbb{E}|g(X_1)| < \infty$. Then*

$$\hat{\mu}_N \xrightarrow{\text{a.s.}} \mu \quad \text{as } N \rightarrow \infty. \quad (1.3)$$

Proof. Set $Y_b = g(X_b)$. Then Y_1, Y_2, \dots are i.i.d. with finite mean $\mu = \mathbb{E}[Y_1]$. Kolmogorov's strong law for i.i.d. integrable sequences [1] states that $N^{-1} \sum_{b=1}^N Y_b \rightarrow \mathbb{E}[Y_1]$ almost surely. The conclusion follows by identifying the partial sums on the left with $\hat{\mu}_N$. \square

Theorem 1.1 is the existence theorem for Monte Carlo. It guarantees that, given enough replicates, the simulator's reported mean converges to the population mean. It says nothing about how *fast* the convergence happens, and it does not on its own quantify the simulator's uncertainty at a finite N . The standard error at finite N is the central limit theorem's contribution.

1.1.2 Central limit theorem and the MCSE bound

Theorem 1.2 (CLT for the Monte Carlo estimator). *Let X_1, X_2, \dots be i.i.d. with distribution F and $\sigma_g^2 = \text{Var } g(X_1) < \infty$. Then*

$$\sqrt{N}(\hat{\mu}_N - \mu) \xrightarrow{d} \mathcal{N}(0, \sigma_g^2) \quad \text{as } N \rightarrow \infty. \quad (1.4)$$

Proof. The summands $Y_b - \mu$ are i.i.d. with mean zero and variance $\sigma_g^2 \in (0, \infty)$. The Lindeberg-Lévy CLT applies. \square

The standard Monte Carlo reference for the standard-error envelope and the practical implications of the CLT is [2]. The CLT delivers the **Monte Carlo standard error** of $\hat{\mu}_N$:

Chapter 3

Mack's Distribution-Free Model

Chapter 2 closed with the residual-bootstrap reserve distribution wrapped around the chain-ladder point estimate \widehat{R} . That envelope cost a distributional commitment — the over-dispersed-Poisson family on the incremental triangle — and a few thousand replicates of CPU time. For the Solvency II SCR calculation the cost is acceptable; for a back-of-envelope second opinion at year-end committee, or for a reserve-risk number that has to be defended without an exponential-family assumption the data may or may not support, the bootstrap is too much machinery for the question being asked.

This chapter delivers the analytic alternative. Mack's 1993 framework keeps the same chain-ladder point estimate \hat{f}_j but retreats the inferential commitment to two moment assumptions — a conditional mean and a conditional variance parameterized by a single development-period volatility σ_j — and derives the mean square error of prediction of the reserve in closed form, no resampling required. The deliverable is the analytic mean-plus-MSE engine, a parametric-bootstrap variant that converts those moments into a full reserve distribution when the regulator's Q(0.995) reading needs one, a multi-line extension for portfolios where each line is fit independently, and the empirical reconciliation against the Chapter 2 bootstrap that closes the analytic-vs-simulation equivalence. Quoting the bootstrap's standard deviation as Mack's analytic prediction error — they are not the same quantity and differ by twenty to forty percent on a typical motor portfolio — is the kind of conflation I have walked into reserving committees ready to refuse. The chapter is built in five sections. §3.1 fixes the two moment assumptions and the unbiased age-to-age sigma estimator. §3.2 derives the closed-form

MSEP recursion and the asymptotic-normality result. §3.3 turns the analytic moments into a full reserve distribution via the parametric bootstrap. §3.4 extends to multi-line portfolios with the Quarg-Mack trend correction and the sigma-floor heuristic. §3.5 closes the chapter with the Mack-vs-BCL bootstrap comparison on `motor_nordic`, `property_continental`, and `multiline_casualty`.

3.1 Mack’s first- and second-moment assumptions and the age-to-age sigma estimator

Let $C_{i,j}$ denote the cumulative paid loss for accident year $i \in \{1, \dots, I\}$ at development period $j \in \{1, \dots, J\}$, with the upper-left triangle observed and the lower-right triangle awaiting projection. Mack [10] (later extended in [15]) introduced a two-assumption framework that makes the chain-ladder estimator inferentially honest without committing to any distribution.

The framework’s central move is the *retreat* from a full likelihood to a moment specification. The Chapter 2 ODP bootstrap treated the chain-ladder point estimate as the MLE under the over-dispersed-Poisson family on the incremental triangle — a distributional commitment that delivered Pearson-residual exchangeability and the full reserve distribution at the cost of a likelihood the data must support. Mack’s framework keeps the chain-ladder point estimate but reads it through a *different* inferential lens: the cumulative paid is modeled as a stochastic process whose conditional first moment is the volume-weighted development factor and whose conditional second moment is parameterized by a single per-column volatility σ_j . The third moment, the parametric family, the residual distribution — none of that is specified. The estimator is what it always was; the inferential machinery underneath is leaner.

The reduced commitment buys robustness. A portfolio whose incremental triangle fails the ODP Pearson-residual heteroscedasticity diagnostic of Chapter 2 §2.2 may still satisfy Mack’s mean-and-variance specification, and the analytic MSEP delivered by §3.2 is then the right inferential quantity. The cost is that the framework no longer produces a *distribution* — only a mean and a variance. §3.3 closes that gap with a parametric bootstrap whose distributional commitment is taken later, at the reporting stage, rather than embedded in the likelihood. The result is a two-stage discipline — analytic moments first, parametric distribution second — that splits the modeling commitments the bootstrap of Chapter 2 made jointly.

Assumption M1 (Conditional mean). *There exist constants f_1, \dots, f_{J-1}*

such that for every accident year i and development period $j \leq J - 1$,

$$\mathbb{E}[C_{i,j+1} | C_{i,1}, \dots, C_{i,j}] = f_j \cdot C_{i,j}. \quad (3.1)$$

Assumption M2 (Conditional variance). *There exist constants $\sigma_1^2, \dots, \sigma_{J-1}^2 > 0$ such that for every accident year i and development period $j \leq J - 1$,*

$$\text{Var}(C_{i,j+1} | C_{i,1}, \dots, C_{i,j}) = \sigma_j^2 \cdot C_{i,j}. \quad (3.2)$$

Assumption M3 (Cross-accident-year independence). *Accident years are independent: $C_{i,\cdot}$ and $C_{k,\cdot}$ are independent for $i \neq k$.*

Assumptions M1 and M2 are weaker than the ODP family of Chapter 2: M2 declares only the *shape* of the conditional variance (proportional to the current cumulative paid), without committing to a likelihood. The bootstrap of Chapter 2 needed the ODP family to define Pearson residuals as exchangeable; Mack's model does not need exchangeability. M1 is the *same* deterministic chain-ladder relation as before, now stated as a conditional mean rather than as a defining formula. M3 — accident-year independence — was implicit in the bootstrap-resampling step of Algorithm 2.8 and is now declared explicitly.

Definition 3.1 (Development factor under Mack's model). *The volume-weighted estimator of the development factor f_j from the observed upper-left triangle is*

$$\hat{f}_j = \frac{\sum_{i=1}^{I-j} C_{i,j+1}}{\sum_{i=1}^{I-j} C_{i,j}}, \quad j = 1, \dots, J - 1. \quad (3.3)$$

The estimator \hat{f}_j is identical to the volume-weighted chain-ladder factor of Definition 2.1. What changes is the inferential interpretation: under M1, \hat{f}_j is the unique best linear unbiased estimator (BLUE) of f_j given the observed column- j and column- $j + 1$ cells, and its conditional variance is what M2 will deliver.

Definition 3.2 (Age-to-age sigma estimator). *The Mack 1993 estimator of the development-period volatility σ_j^2 under M2 is*

$$\hat{\sigma}_j^2 = \frac{1}{I-j-1} \sum_{i=1}^{I-j} C_{i,j} \left(\frac{C_{i,j+1}}{C_{i,j}} - \hat{f}_j \right)^2, \quad j = 1, \dots, J - 2. \quad (3.4)$$

The estimator is unbiased for σ_j^2 under M1, M2, M3 — the proof is a direct computation that the column- j cumulative paid acts as the exposure weight

Monte Carlo Methods for Non-Life Actuaries delivers the full reserving-and-capital pipeline a Solvency II internal-model team needs but no single text on the actuary's shelf assembles. Glasserman speaks to derivatives books. Mack is the analytic core but is silent on cat, copula, and allocation. Klugman covers compound aggregates but stops short of the one-year horizon. The non-life carrier stitches them together from scratch every recalibration cycle.

Dr. Sigrid Norrgård — twenty-six years across If P&C reserving and Munich Re Stockholm capital modelling, sign-off actuary on roughly one hundred and forty year-end reserve reviews and twelve internal-model approval submissions — provides that stitch. Fifteen chapters cover the methods (bootstrap chain ladder, Mack analytic, Tweedie GLM, Bayesian MCMC, Panjer compound, GPD severity tail, t-copula aggregation, reinsurance layer pricing, ELT-driven cat, Merz–Wüthrich one-year, premium-risk cycle, Euler and Aumann–Shapley allocation) on real four-line calibrated portfolios. Every method ships with its R implementation — a working production stack you can run, fork, and walk into a year-end review with. The hundred-and-five-exercise appendix is the methodology defense in the actuary's voice: the questions a reviewer will ask, answered.

For the actuary whose next internal-model submission must read defensibly to the Finansinspektionen, BaFin, or the PRA — and who has stopped pretending the existing literature covers what the regulator actually asks.

