

Identification and Bayesian Inference of Loss-Development Structure

Model-Averaged Breakpoints in the Development and Calendar-Year Effects,
with an Application to a Workers' Compensation Paid-Loss Triangle

Frank Schmid

This version: 3 July 2026
First version: 7 June 2026

Contents

Abbreviations and Acronyms	iii
Greek Letters	iii
Executive Summary	iv
1 Background and Related Literature	1
2 The Breakpoint Model	2
2.1 Multiplicative structure and the cumulative observation	2
2.2 The age–period–cohort identification constraint	3
2.3 Calendar identification by anchoring to an external macroeconomic level	3
2.4 The development-year effect as growth-rate regimes	3
2.5 Calendar forecast: the Ornstein–Uhlenbeck deviation	4
2.6 The accident-year levels as a first-order random walk	4
2.7 Priors	5
2.8 Inference: collapsed reversible jump and exact enumeration	5
2.9 Predictive distribution by simulation	6
3 Data and Scope	7
4 Results	9
4.1 Number and location of breakpoints	9
4.2 Leading segmentations and growth-rate regimes	11
4.3 Development-year effect and shares	13
4.4 Accident-year effect	14
4.5 Calendar-year effect	15
4.6 Calendar forecast: the anchored backbone and a reverting deviation	16
4.7 The early-1990s calendar departure: a closed breakpoint	18
4.8 The value of the calendar-year anchor	21
4.9 Forward calendar uncertainty	23

5	Diagnostics and Validation	23
5.1	Goodness of fit	24
5.2	Fitted versus observed cumulative	24
5.3	Log-deviation triangle	25
5.4	Residuals by accident, development, and calendar direction	26
5.5	Collapsed reversible jump against exact enumeration	28
5.6	Sensitivity to the rate prior	29
6	Loss Reserve Estimate and Predictive Distribution	29
6.1	Completed projection	29
6.2	Posterior of the total outstanding loss	30
6.3	Sensitivity of the loss reserve estimate to a calendar-year shock	31
6.4	The 1997 ultimate and the structural component	31
7	Modeling Considerations	34
7.1	The heteroskedastic cumulative observation model	34
7.2	A misspecification-robust posterior for the cumulation correlation	35
7.3	Inference for the latest accident year	36
8	Conclusion	37
A	Paid-Loss Triangles	38
B	Finite-Variance Reporting of the Predictive Distribution	39
B.1	Log-scale dispersion and quantiles	39
B.2	A thin-tailed prior for the observation variance	39
C	Evidence Profile of the Deviation Parameters	41
	Bibliography	43

Abbreviations and Acronyms

APC	Age–period–cohort
AY	Accident year
BF	Bayes factor
CAS	Casualty Actuarial Society
DY	Development year
GDP	Gross domestic product
GIG	Generalized inverse Gaussian
IMF	International Monetary Fund
LOWESS	Locally weighted scatterplot smoothing
MCMC	Markov chain Monte Carlo
NAIC	National Association of Insurance Commissioners
NCCI	National Council on Compensation Insurance
NGDP	Nominal gross domestic product
OU	Ornstein–Uhlenbeck
RJMCMC	Reversible-jump Markov chain Monte Carlo
RMSE	Root-mean-square error
SD	Standard deviation
TTD	Temporary total disability
USD	United States dollar
WEO	World Economic Outlook (IMF)

Greek Letters

α (alpha)	Accident-year effect
β (beta)	Development-year effect
γ (gamma)	Calendar-year effect
δ (delta)	Trend increment in the identification argument
Δ (delta)	Per-year change in the development log-effect (growth rate)
ε (epsilon)	Error term
η (eta)	Prior mean of the truncated-Poisson breakpoint count
θ (theta)	Ornstein–Uhlenbeck reversion rate
μ (mu)	μ_0 , the intercept of the log-incremental mean
ξ (xi)	Random-walk innovation of the accident-year levels
ρ (rho)	Posterior-weighted evidence information ratio
σ (sigma)	Standard deviation
τ (tau)	Random-walk innovation standard deviation of the accident-year levels
ϕ (phi)	Ornstein–Uhlenbeck persistence, $\phi = e^{-\theta}$
χ (chi)	Chi, in the scaled inverse-chi-squared full conditional of the observation variance
ψ (psi)	Psi, the right-tail parameter of the generalized inverse Gaussian prior on the observation variance

Executive Summary

The structure of a loss-development model—the shape of the run-off across development years, the relative levels of the accident-year cohorts, and the calendar-year trend—is, in standard reserving practice, specified as part of the model, and the outstanding loss is estimated conditional on that specification. This paper takes the structure as an object of inference. It places the development-year effect, the accident-year levels, and the calendar-year effect under a single Bayesian model, infers them jointly from the triangle, and carries the resulting uncertainty about the structure through to the loss reserve estimate and to the predictive distribution of the most recent accident-year ultimate.

The development-year effect is represented as a sequence of growth-rate regimes: the per-year change in the development log-effect is constant within a regime and changes only at a small number of breakpoints, whose number and location are unknown and are inferred by reversible-jump Markov chain Monte Carlo. Because the development axis admits a candidate set small enough to enumerate, the posterior over all segmentations is also computed exactly, and the exact posterior serves as an independent check that the sampler reproduces. The accident-year levels carry a first-order random-walk prior that pools information across adjacent cohorts. The calendar-year effect, which the age–period–cohort restriction leaves unidentified along one linear direction, is identified by anchoring it to the level of an external macroeconomic series and, on the same model-averaged footing as the development-year regimes, admits a localized break whose presence and location are inferred. The cumulative paid amount is observed log-normally about a multiplicative mean. The outstanding loss and its predictive distribution follow from the posterior by simulation, the predictive uncertainty summarized by its analytic mean, median, 90 percent interval, and log-scale dispersion—the linear-scale variance being sensitive to a heavy upper tail—and separated into an observation component and a point-forecast component, the latter carrying parameter, structural, and calendar-forecast uncertainty.

The method is illustrated throughout on a single paid-loss triangle, a State Farm workers' compensation triangle, the line on which the calendar-year effect is traditionally most pronounced and on which the nominal-gross-domestic-product anchor is economically grounded through an inflation channel and a real-activity channel together. The illustration recovers an interpretable run-off, described by a single development-year breakpoint, and a calendar-year trend anchored to the level of nominal gross domestic product, its linear slope fixed to the anchored backbone, its curvature a mean-reverting deviation that the triangle identifies, and an early-1990s calendar-year departure carried by a closed breakpoint as the lagged effect of the 1990–91 recession, a representation to which the total loss reserve estimate is robust. The inferred effects, the loss reserve estimate, and the predictive distribution of the ultimate for that triangle are reported in the body.

1 Background and Related Literature

Stochastic loss reserving treats the run-off triangle as the realization of a statistical model and reports the loss reserve estimate together with a measure of its uncertainty. A broad class of these models decomposes the triangle along three directions—the accident year, the development year, and the calendar year—and writes the incremental or cumulative payment as a function of an accident-year (cohort) effect, a development-year (run-off) effect, and a calendar-year (diagonal) effect. England and Verrall survey the stochastic-reserving models that descend from the chain-ladder technique, among them the parametric curves that reduce the development-year parameters and allow the run-off to be extrapolated and the Bayesian formulations that deliver a full predictive distribution of the outstanding loss.¹ The distribution-free standard error of the chain-ladder reserve estimate is due to Mack,² and the Bayesian, Markov-chain-Monte-Carlo strand of the reserving literature is developed at length by Meyers.³ The calendar-year direction carries the trend in settlement cost—claims inflation and superimposed inflation—and is the third axis of the probabilistic-trend formulation of Barnett and Zehnwrith, in which a mis-specified calendar-year trend distorts the loss reserve estimate.⁴

The three directions are not free of one another. Because the calendar year is the sum of the accident year and the development year, a linear trend can be shifted among the three axes without altering any fitted value: the level and the curvature of each effect are identified, but the three linear slopes are not separately estimable. This is the age–period–cohort identification problem, characterized for the extended chain-ladder model by Kuang, Nielsen, and Nielsen.⁵ Resolving it requires one external restriction. Schmid introduced the two devices on which the present paper rests: he anchored the calendar-year effect to an external economic rate—a consumption, or inflation, path—fixing the otherwise unidentified linear direction, and he represented the development-year effect as a piecewise-linear function whose breakpoints are not fixed in advance but are inferred, with model averaging, by reversible jump. That treatment was applied to workers’-compensation indemnity and medical triangles, the line the present paper treats.⁶ Reversible-jump Markov chain Monte Carlo has also been applied to claims reserving by Verrall and Wüthrich, there to reduce the parameter count of the chain-ladder model and to estimate the tail;⁷ the present application is structural rather than dimensional, inferring the number and the location of slope changes in the development-year and calendar-year effects.

This paper adopts both devices. The development-year effect is represented as a sequence of growth-rate regimes, a piecewise-linear development log-effect whose breakpoints—their number and their location—are inferred by reversible jump and, because the development axis admits a candidate set small enough to enumerate, are additionally checked against the exact posterior obtained by enumeration. The calendar-year effect is anchored to the level of an external macroe-

¹Peter D. England and Richard J. Verrall, “Stochastic Claims Reserving in General Insurance,” *British Actuarial Journal* 8, no. 3 (2002): 443–518.

²Thomas Mack, “Distribution-Free Calculation of the Standard Error of Chain Ladder Reserve Estimates,” *ASTIN Bulletin* 23, no. 2 (1993): 213–225.

³Glenn G. Meyers, *Stochastic Loss Reserving Using Bayesian MCMC Models*, CAS Monograph Series 1 (Arlington, VA: Casualty Actuarial Society, 2015).

⁴Glen Barnett and Ben Zehnwrith, “Best Estimates for Reserves,” *Proceedings of the Casualty Actuarial Society* 87 (2000): 245–321.

⁵D. Kuang, B. Nielsen, and J. P. Nielsen, “Identification of the Age–Period–Cohort Model and the Extended Chain-Ladder Model,” *Biometrika* 95, no. 4 (2008): 979–986.

⁶Frank A. Schmid, “The Workers Compensation Tails,” *Variance* 6, no. 1 (2012): 48–77.

⁷Richard J. Verrall and Mario V. Wüthrich, “Reversible Jump Markov Chain Monte Carlo Method for Parameter Reduction in Claims Reserving,” *North American Actuarial Journal* 16, no. 2 (2012): 240–259.

conomic series, the linear trend fixed to the anchored backbone and the curvature carried by a bounded mean-reverting deviation; the accident-year levels carry a first-order random-walk prior that pools information across adjacent cohorts. The remainder of the paper is organized as follows. Section 2 sets out the breakpoint model, its identification, and the inference; Section 3 describes the triangle and the macroeconomic series; Sections 4 and 5 report the inferred structure and the diagnostics; Section 6 gives the loss reserve estimate and the predictive distribution; Section 7 discusses the observation model; and Section 8 concludes.

2 The Breakpoint Model

The incremental payment carries three additive effects on the log scale: an accident-year effect, a development-year effect, and a calendar-year effect. The development-year effect is the object of the analysis and is modeled as a small number of growth-rate regimes, and the calendar-year effect is not separately identified and is fixed by reference to an external macroeconomic rate. The accident-year levels carry a first-order random-walk prior. The subsections set out the structure and identification, the regime construction, the calendar forecast, the priors, the inference by reversible jump and exact enumeration, and the predictive simulation.

2.1 Multiplicative structure and the cumulative observation

Let $i = 0, \dots, 9$ index the accident year and $j = 0, \dots, 9$ the development year. The incremental paid amount has a multiplicative mean in an accident-year effect, a development-year effect, and a calendar-year effect,

$$m_{ij} = \exp(\mu_0 + \alpha_i + \beta_j + \gamma_t), \quad t = i + j, \quad (1)$$

where μ_0 is an intercept, α_i are the accident-year (cohort) levels, β_j are the development-year effects, and γ_t are the calendar-year (diagonal) effects, with $\alpha_0 = 0$ and $\beta_0 = 0$ for identification. The cumulative amount is the partial sum across development,

$$M_{ij} = \sum_{k=0}^j m_{ik} = \exp(\mu_0 + \alpha_i) \sum_{k=0}^j \exp(\beta_k + \gamma_{i+k}), \quad (2)$$

and the observed cumulative amount is log-normal about this mean,

$$\log C_{ij} = \log M_{ij} + \varepsilon_{ij}, \quad \varepsilon_{ij} \sim \mathcal{N}(0, \sigma^2 v_{ij}), \quad (3)$$

where the cell-variance factor v_{ij} carries a development-year-varying observation scale, constructed in Section 7.1 from a log-quadratic incremental scale by cumulation-aware accumulation; the homoskedastic special case is $v_{ij} \equiv 1$. Working with the cumulative amount on the log scale keeps the observation model defined when an incremental amount is small or negative, and equation (2) shows that the accident-year level $\mu_0 + \alpha_i$ factors out of the cumulative sum. The accident-year level is therefore a log-additive term, a property used in estimation. The mean of equation (2) is a nonlinear (log-sum-exp) function of the effects. Within an accident year the cumulative cells are partial sums of common increments and so are positively correlated; the observation model weights them as independent, a diagonal working likelihood whose set-aside correlation is carried after estimation by the misspecification-robust posterior of Section 7.2, which enters both the comparison of segmentations and the predictive interval.

2.2 The age–period–cohort identification constraint

Because $t = i + j$, the three axes carry an exact linear dependency, the age–period–cohort (APC) identification problem. For any constant δ , the substitution

$$\alpha_i \rightarrow \alpha_i - \delta i, \quad \beta_j \rightarrow \beta_j - \delta j, \quad \gamma_t \rightarrow \gamma_t + \delta t \quad (4)$$

leaves each incremental mean in equation (1) unchanged, and therefore leaves the cumulative mean of equation (2) and every fitted value unchanged. A linear trend can be moved among the three axes without altering the fit. The level, the curvatures (second differences) of each axis, and the within-sample fitted values are identified; the individual linear slopes are not.⁸ The aliasing extends to the projection: a future cell at calendar year $t = i + j$ carries the same accident-year level and development-year effect as the slide leaves them, so the substitution leaves every projected incremental mean, and therefore the outstanding loss and the ultimate, unchanged as well. A constant added to the baseline calendar-year rate of growth is re-absorbed by the development-year and accident-year effects in sample and in the projection alike, so the loss reserve estimate does not move with the level of the calendar-year anchor. What the inference draws from the anchor is its shape—the curvature it contributes and the long-run trend it sets—and not its level.

2.3 Calendar identification by anchoring to an external macroeconomic level

The free dimension is fixed by one linear restriction. The calendar-year effect is anchored to the logarithm of an external macroeconomic level: its linear trend is set so that the average year-over-year increment over the observed diagonals equals the average logarithmic rate of growth of the external series, realized over the observed window (Section 3). Under the APC slide of equation (4) the average increment changes by δ , so the restriction determines δ uniquely. The calendar-year effect is the anchored backbone plus a mean-reverting deviation,

$$\gamma_t = c_t + u_t, \quad (5)$$

where c_t is the backbone, the cumulative logarithm of the external level, expected to track the calendar-year effect up to an additive constant, and u_t is an Ornstein–Uhlenbeck (OU) deviation carried on the curvature space orthogonal to the constant and the linear trend. The anchoring fixes the linear trend of the calendar-year effect to the average slope of the backbone, while the curvature is the deviation, mean-zero over the observed diagonals with no identifiable in-sample spread. This tracking is approximate and holds over the long run: the anchoring equates the average logarithmic rates of growth of the calendar-year effect and the external level, while their shorter-run movements remain unrelated. The relabeling leaves every fitted value unchanged.

2.4 The development-year effect as growth-rate regimes

The development log-effect β_j is the cumulative sum of per-year growth rates. Define the growth rate at development year j as $\Delta\beta_j = \beta_j - \beta_{j-1}$. In the breakpoint model the growth rate is piecewise constant: it changes only at a set $B \subseteq \{2, 3, \dots, 9\}$ of breakpoints. With $K = |B|$ breakpoints the development axis is partitioned into $K + 1$ segments, each carrying a single rate. Writing $g(j)$ for the segment containing development year j and $r_0, \dots, r_K$ for the segment rates,

$$\beta_0 = 0, \quad \beta_j = \beta_{j-1} + r_{g(j)} \quad (j = 1, \dots, 9). \quad (6)$$

⁸Kuang, Nielsen, and Nielsen, “Identification,” 979–986.

The development log-effect is thus continuous and piecewise linear in j : the level is continuous at a breakpoint, and the slope changes. A breakpoint marks a development year at which the rate of run-off shifts from one constant value to another. The candidate breakpoint positions are the six development years $\{2, \dots, 7\}$, giving $2^6 = 64$ possible segmentations. A candidate is admitted only where the post-break regime it creates rests on at least three increments—three slope-defining differences—so that the post-break rate carries a degree of freedom beyond the exact identification that two increments would leave. The rate change at development year j is carried by the increments $\Delta\beta_j$ through $\Delta\beta_9$, so three or more requires j at or before 7, and the candidate set stops at development year 7. A breakpoint at development year 8 or 9 would leave the post-break rate on two increments or on one, resting on the sparsest cells of the triangle, which at development year 9 number one. The earliest candidate is development year 2, since a breakpoint at development year 1 would leave the first segment a single year and its rate unidentified. The same three-increment criterion governs the calendar axis in Section 4.7.

2.5 Calendar forecast: the Ornstein–Uhlenbeck deviation

The future diagonals lie beyond the latest observed diagonal and require calendar-year effects that are not estimated within the triangle. Under the anchoring of Section 2.3 the calendar-year effect is the backbone plus the deviation, so the forecast is supplied by the projected backbone and the reverting deviation. The backbone is read forward from the projected level of the external series; the deviation u_t follows an Ornstein–Uhlenbeck process, mean-reverting to zero,

$$u_{k+1} = \phi u_k + \varepsilon_k, \quad \phi = e^{-\theta}, \quad \text{Var}(\varepsilon_k) = V(1 - \phi^2), \quad (7)$$

the discrete annual form of $du = -\theta u dt + \sigma dW$, with stationary variance $V = \sigma^2/(2\theta)$ and reversion rate θ . The reversion rate is tested for identification against the calendar curvature the triangle resolves; where the data do not identify it, the deviation reduces to its one-parameter form, independent innovations ($\phi = 0$) under the single stationary variance V . Forward, the deviation reverts to zero at rate ϕ , so the central calendar path is the projected backbone and the forecast variance is bounded by V rather than growing without limit. The anchor is annual, so each calendar year carries a single level. The projection covers six future diagonals, beyond which the backbone is extended at the last observed annual growth rate; the projected path is reported in Section 4.6. A future incremental payment in a given calendar year carries the calendar-year effect of that year, so the forecast applies to the payment of each year and not to amounts already paid.

2.6 The accident-year levels as a first-order random walk

The accident-year levels are given a first-order random-walk prior on the log scale,

$$\alpha_i = \alpha_{i-1} + \xi_i, \quad \xi_i \sim \mathcal{N}(0, \tau^2), \quad i = 1, \dots, 9, \quad (8)$$

with $\alpha_0 = 0$ and a flat prior on the overall level. The prior couples adjacent cohorts, so a cohort with few observed cells is informed by its neighbors. The innovation standard deviation τ governs the strength of the coupling: a small τ pulls the levels toward a common path, a large τ frees them. It carries a half-normal prior with scale 0.5 and is estimated jointly with the remaining parameters. The prior matters most for the most recent accident year, which is observed only at development year 0: its level is identified through the random walk from the adjacent cohorts rather than from its single cell, a point taken up in Section 7.3.

2.7 Priors

The number of breakpoints carries a truncated-Poisson prior with mean $\eta = 1$, which favors a parsimonious run-off structure. Given K , the breakpoint positions are uniform over the $\binom{6}{K}$ subsets of size K of the six candidate development years (Section 2.8). The segment rates carry independent normal priors with a common variance, $r_g \sim \mathcal{N}(0, \sigma_r^2)$, and the variance carries an inverse-gamma hyperprior that is integrated in closed form, so the segment rates acquire a multivariate- t marginal whose scale the data set. Integrating the rate variance rather than fixing it carries the uncertainty in the rate scale into the comparison of segmentations and into the rate posterior, and does not condition the inference on one value of a scale the data identify weakly: the marginal likelihood of the development segmentation is nearly flat in σ_r over the range the data support. The hyperprior is centered at the scale of the triangle's own development rates rather than at a constant carried from another triangle. That centering scale is the root-mean-square of the saturated free-rate development slopes—a separate growth rate at every development year, at its maximum-likelihood value—taken over the development years that at least four accident years observe; for the present triangle the scale is $\sigma_r = 0.5782$. The four-observation floor excludes the late development years, which rest on one to three cells and whose free rates do not stabilize; it is the floor the candidate breakpoint set also enforces, and the centering scale is stable with respect to the width of the integrating prior. The calendar-year effect carries the Ornstein–Uhlenbeck deviation prior of Section 2.5, a stationary mean-reverting prior on the curvature space orthogonal to the constant and the linear trend, with weakly-informative half-life and amplitude and its reversion rate tested for identification (both parameters are audited by marginal likelihood in Appendix C); the accident-year levels carry the first-order random-walk prior of Section 2.6, whose innovation standard deviation τ is itself estimated under a half-normal prior with scale 0.5. Because the marginal likelihood of a segmentation integrates the segment rates against this prior, the rate scale enters the comparison of segmentations through the parsimony penalty; the sensitivity to the scale is reported in Section 5.6.

The breakpoint-count rate is set to be scale-aware, so that the method carries to longer triangles than the ten-year triangle treated here. For the present triangle the rate is $\eta = 1$, the value at which the prior expects one breakpoint among the six candidate development years. As the development axis lengthens, the candidate count m grows, and the rate grows with it, $\eta(m) = \max(1, 1 + \log(m/m_0))$, anchored at the standard ten-year triangle, $m_0 = 6$, so the prior expects one breakpoint for the standard triangle and proportionally more, sub-linearly, for a longer triangle—about two breakpoints for a twenty-year triangle and about three for a forty-period triangle. The uniform-placement term $-\log \binom{m}{K}$ retains the automatic multiplicity adjustment as m grows, tightening the prior odds of any single breakpoint as the candidate set widens.⁹ At $m = 6$ the rate equals one and the scale-aware form reproduces the fixed prior, so it does not affect the inference reported here.

2.8 Inference: collapsed reversible jump and exact enumeration

The number and location of breakpoints, the segment rates, and the remaining parameters are inferred by reversible-jump Markov chain Monte Carlo (RJMCMC).¹⁰ Conditional on a segmentation the model is, in the continuous parameters, a penalized Gaussian regression: the accident-year

⁹The placement term is the automatic multiplicity adjustment of James G. Scott and James O. Berger, “Bayes and Empirical-Bayes Multiplicity Adjustment in the Variable-Selection Problem,” *Annals of Statistics* 38, no. 5 (2010): 2587–2619.

¹⁰Peter J. Green, “Reversible Jump Markov chain Monte Carlo Computation and Bayesian Model Determination,” *Biometrika* 82, no. 4 (1995): 711–732.

levels, the segment rates, and the calendar curvature enter through the cumulative observation of equation (2), and the observation variance and the random-walk innovation standard deviation carry conjugate priors. These continuous parameters are integrated out by the Laplace approximation, the observation variance concentrated out and the innovation standard deviation integrated over a grid, leaving the marginal evidence of the segmentation. The trans-dimensional chain moves over the marginal posterior of the segmentation alone, accepting a proposed change of the breakpoint set on the ratio of marginal evidences and segmentation priors. Collapsing the continuous parameters in this way removes the matching of parameters across dimensions and the unit-Jacobian auxiliary draw that a joint birth-and-death move carries, which is the slower-mixing component of the joint sampler (Section 5.5).

The segmentation proposals are data-informed: a birth weights the candidate locations by the marginal evidence each would attain, and the matching death and a relocation move are balanced against it, so the chain proposes changes where the evidence supports them.¹¹ Conditional on the sampled segmentation, the continuous parameters are drawn from their full conditionals: the accident-year levels jointly from their Gaussian conditional under the random-walk prior of equation (8); the innovation standard deviation by a griddy-Gibbs step under its half-normal prior; the observation variance from its scaled inverse-chi-square conditional; and the segment rates and the calendar curvature by random-walk Metropolis.

Because the development axis admits only six candidate positions, the marginal posterior over the 64 segmentations is also computed exactly, segmentation by segmentation, by the same Laplace evidence combined with the breakpoint-count and location priors.¹² The collapsed chain is validated against this enumeration in Section 5.5; the enumeration is the exactness benchmark where the candidate set is small enough to enumerate, and the collapsed chain is the general method where it is not. Reported conditional and predictive quantities are Rao-Blackwellized over the segmentations: each segmentation contributes to the outstanding loss, the ultimate, and the calendar-year trend in proportion to its posterior probability, with independent within-segmentation draws of the continuous parameters and the calendar forecast.¹³ Because the segmentation weighting is exact and the within-segmentation draws are independent, the reported predictive carries no between-segmentation autocorrelation and its effective sample size equals the simulation size.

2.9 Predictive distribution by simulation

The predictive distribution of an ultimate is obtained by Monte Carlo from four sources. The first is observation (process) noise on the projected ultimate. The second is parameter uncertainty in the segment rates, the accident-year levels, and the calendar curvature. The third is structural uncertainty: the segmentation is sampled from its posterior, so the run-off structure itself varies across draws. The fourth is calendar-forecast uncertainty: the Ornstein–Uhlenbeck deviation is drawn from its stationary law and reverts to the projected backbone, so the future calendar path is simulated under the bounded deviation variance. The structural and parameter sources enter together through the Rao-Blackwellized mixture over segmentations; the law-of-total-variance separation of the structural component is reported in Section 6.4.

¹¹Giacomo Zanella, “Informed Proposals for Local MCMC in Discrete Spaces,” *Journal of the American Statistical Association* 115, no. 530 (2020): 852–865.

¹²Robert E. Kass and Adrian E. Raftery, “Bayes Factors,” *Journal of the American Statistical Association* 90, no. 430 (1995): 773–795.

¹³Jennifer A. Hoeting, David Madigan, Adrian E. Raftery, and Chris T. Volinsky, “Bayesian Model Averaging: A Tutorial,” *Statistical Science* 14, no. 4 (1999): 382–417.

3 Data and Scope

The model of Section 2 is applied to a paid-loss triangle of ten accident years (AY) from 1988 through 1997 by ten development years (DY) from 0 through 9, in thousands of United States dollars (USD).¹⁴ The triangle holds 55 observed cells on and above the latest diagonal; the 45 cells below it are unobserved and are the object of projection. The cumulative payments are formed as the running sums of the incremental source integers, so no source integer is rounded, and every cumulative value is positive. The incremental and cumulative triangles are reproduced in Appendix A (Tables A.1 and A.2). The macroeconomic input is United States nominal gross domestic product (NGDP), annual, from the International Monetary Fund (IMF) World Economic Outlook (WEO). Over the observed window, nominal gross domestic product grew at an average of 5.4891 percent per year on the log scale; this realized average is the average slope of the calendar-year backbone c_t of equation (5).

The choice of nominal gross domestic product as the calendar-year anchor rests on the economics of the workers' compensation calendar-year effect. Workers' compensation is the line on which the calendar-year, or diagonal, effect is traditionally most pronounced, and its movement is tied to economic conditions through two channels, an inflation channel and a real-activity channel, which nominal gross domestic product carries together as the sum of inflation and real growth.

The inflation channel is medical-price inflation. The indemnity, or wage-replacement, portion of a continuing claim is subject to statutory indexation that varies across jurisdictions—at a fixed rate in some, tied to a price or wage index in others, and absent in others—while the medical portion is subject to medical-price inflation. Medical prices in workers' compensation are in large part administered prices, set by fee schedules, and workers'-compensation physician-price inflation tends to run below the medical-care and physician-services components of the consumer price index.¹⁵ Because medical prices are administered, medical inflation is less responsive to variation in economic activity than general inflation; the price channel feeds the long-term calendar trend and contributes little to its short-term cyclical variation.

The real-activity channel has two sub-channels of opposite sign and different horizon. The first is a duration channel that operates over the business cycle. An economic slowdown lengthens the duration of continuing claims across the open diagonals: a weaker labor market lowers the propensity to return to work for claims of every accident-year age at once, so the calendar-year effect moves opposite to real growth over the cycle. Temporary total disability (TTD) duration tracks the unemployment rate, lengthening in the recession and shortening in the recovery, and the lengthening appears across several accident years at once on the same calendar-year diagonals, which is the calendar signature of the effect.¹⁶ The second is a utilization channel that operates over the long run: growth in real gross domestic product per capita raises the per-capita consumption of medical services over time, a secular relationship rather than a contemporaneous one. The medical

¹⁴The triangle is the net paid-loss triangle for the Workers' Compensation line of State Farm Mutual Automobile Insurance Company (National Association of Insurance Commissioners group code 1767), drawn from the Schedule P loss-reserving database compiled by Meyers and Shi from annual-statement filings. Glenn G. Meyers and Peng Shi, "The Retrospective Testing of Stochastic Loss Reserve Models," *Casualty Actuarial Society E-Forum* (Summer 2011): 1–37.

¹⁵Frank Schmid and Nathan Lord, "The Impact of Physician Fee Schedule Changes in Workers Compensation: Evidence From 31 States" (Boca Raton, FL: National Council on Compensation Insurance, 2013).

¹⁶Barry Lipton, John Robertson, and Katy Porter, "Workers Compensation Temporary Total Disability Indemnity Benefit Duration—2013 Update" (Boca Raton, FL: National Council on Compensation Insurance, 2013), Figure 2 of the source. Their Figure 2 shows the recession lifting emerged duration across accident years 2008 through 2010 on the same diagonals.

severity that enters the calendar-year effect is the product of price and utilization, so the utilization channel adds to the long-term trend a component the price channel alone does not carry.¹⁷

Nominal gross domestic product carries both channels. Its inflation component carries the medical-price trend, and its real-growth component carries the secular utilization trend. The utilization channel operates through real growth per capita; total nominal gross domestic product exceeds the per-capita series by the rate of population growth, which is nearly constant over the period, so the two differ by a nearly-constant calendar-year growth rate that the age–period–cohort restriction of Section 2.2 reabsorbs. Total nominal gross domestic product therefore serves as the composite anchor in place of the per-capita series. A medical-price index would mis-measure even the price sub-component, since workers’-compensation price diverges from the medical consumer price index, and would miss the utilization channel and the entire labor-market duration channel, because administered medical prices do not respond to economic activity while claim duration does. Nominal gross domestic product growth summarizes economic conditions, but the force in economic activity most relevant to the calendar-year effect, the variation in claim duration, lags real gross domestic product growth and is not summarized by it. This lagged duration response, together with the workers’-compensation-specific medical inflation, is carried by the curvature of the calendar-year effect about the backbone. The 1990–91 recession enters the calendar-year effect through this lag, on the 1993 and 1994 diagonals, where a closed calendar-year breakpoint carries it (Section 4.7).

Two further considerations support the anchor. The calendar direction of a paid-loss triangle carries the trend in settlement cost, and a mis-specified calendar-year trend distorts the loss reserve estimate.¹⁸ Nominal gross domestic product is a single published series carrying a documented forward projection, so the calendar-year trend is read from one documented source both within the sample and forward, rather than constructed from separate price and real components or extrapolated from in-sample growth. The series is United States nominal gross domestic product, annual.¹⁹ The in-sample backbone is the cumulative logarithm of the realized series over the observed window; the forward backbone is the published projection over the six future diagonals 1998 through 2003, extended over the remaining diagonals 2004 through 2006 at the last observed annual growth rate, read forward as the central calendar path of Section 2.5.

The latest observed cumulative paid amount across the ten accident years sums to USD 1,434,790 thousand. The first-development-year cumulative amounts range from USD 22,190 thousand for accident year 1988 to USD 46,447 thousand for accident year 1993. The 1997 accident year is observed only at development year 0, at USD 25,265 thousand, so its run-off is projected from the estimated structure. The latest observed diagonal is the 1997 calendar year; developed to ultimate, payments extend to 2006.

¹⁷Schmid and Lord, “Physician Fee Schedule Changes,” decompose medical severity growth into a price component and a utilization component.

¹⁸Barnett and Zehnworth, “Best Estimates for Reserves.”

¹⁹International Monetary Fund, *World Economic Outlook Database*, October 2013, United States gross domestic product at current prices in national currency, series NGDP. The realized series is held to the books-close information set; the forward path is treated as extending over six future diagonals, 1998 through 2003, with the remaining diagonals 2004 through 2006 extended at the last observed annual growth rate, the steady-state rate at which the projection terminates.

4 Results

This section reports the inferred run-off structure—the number and location of the breakpoints and the implied growth-rate regimes—followed by the development-year, accident-year, and calendar-year effects and the Ornstein–Uhlenbeck calendar forecast.

4.1 Number and location of breakpoints

Table 1 and Figure 1 give the posterior over the number of breakpoints. One breakpoint carries posterior probability 0.682, two carry 0.277, and three carry 0.037; the single-rate structure and four or more breakpoints carry negligible mass. The run-off is described by one change in the development growth rate.

Table 1: Posterior distribution of the number of breakpoints K , from exact enumeration and from the collapsed reversible-jump chain (four pooled runs). Probabilities below 0.001 are shown as 0.000. The candidate set admits up to six breakpoints; mass beyond three breakpoints is negligible.

K	Exact	Collapsed RJMCMC
1	0.682	0.682
2	0.277	0.278
3	0.037	0.037
4	0.004	0.003




Figure 1: Posterior probability of the number of breakpoints K in the development growth rate. Blue bars are the exact posterior from enumeration of all 64 segmentations with Laplace within-model evidence; orange bars are the collapsed reversible-jump estimate from four pooled runs. The horizontal axis is the breakpoint count; the vertical axis is posterior probability. The segment-rate variance carries an inverse-gamma hyperprior integrated in closed form, centered data-equivariantly at $\sigma_r = 0.5782$, and K carries a truncated-Poisson(1) prior.

Table 2 and Figure 2 give the posterior probability of a breakpoint at each development year. A breakpoint at development year 2 is present in every segmentation. A second breakpoint, where one is present, is weakly located: no candidate development year beyond development year 2 carries inclusion probability above 0.12, and the second breakpoint, when it appears, falls toward the later development years rather than at a single resolved position.

Table 2: Posterior probability of a breakpoint at each development year (inclusion probability), from exact enumeration and from the collapsed reversible-jump chain. A breakpoint at a development year marks a change in the per-year growth rate beginning at that year.

Development year	2	3	4	5	6	7
Exact	1.000	0.046	0.051	0.064	0.085	0.117
Collapsed RJMCMC	1.000	0.046	0.051	0.064	0.085	0.117




Figure 2: Posterior probability that the per-year development growth rate changes at each candidate development year. Blue bars are exact; orange bars are the collapsed reversible-jump estimate. The horizontal axis is the development year of the candidate breakpoint; the vertical axis is posterior inclusion probability.

4.2 Leading segmentations and growth-rate regimes

Table 3 lists the segmentations with the largest posterior probability. The leading segmentation places a single breakpoint at development year 2, with posterior probability 0.682; the next places a second breakpoint at development year 7, with probability 0.091. The leading segmentation describes a two-phase run-off: a positive first-phase rate near $+0.44$ in the development log-effect from development year 0 to development year 1, and a declining tail rate near -0.55 per year from development year 2 onward.

The exact posterior of Table 3 forms each segmentation's evidence from the cumulative cells treated as conditionally independent. Under the heteroskedastic observation model the cumulative cells along an accident year share their early increments, so the diagonal evidence is a composite likelihood rather than the full likelihood. The posterior-weighted information ratio ρ —the trace of $H^{-1}J$ under the independence working evidence relative to its value under the cumulation-correlated evidence, the evidence analog of the misspecification-robust posterior of Section 7.2—is 0.29, so the diagonal evidence over-states the information that distinguishes one segmentation from another. Tempering the within-model evidence by ρ flattens the structure posterior: the single-breakpoint probability, the leading segmentation $\{2\}$, falls from 0.682 to 0.479, the two-breakpoint probability rises from 0.278 to 0.366, and the three-breakpoint probability rises from 0.037 to 0.121, while the certainty of a breakpoint at development year 2 holds. The tempering is reported as a sensitiv-

ity and is not applied as a correction: a single-factor temper is a diagnostic rather than a calibrated adjustment, and the loss reserve is insensitive to the structure weights because the leading segmentations imply a similar run-off (Section 6.4).

Table 3: Segmentations with the largest posterior probability, from exact enumeration. The breakpoint set lists the development years at which the per-year development growth rate changes; K is the number of breakpoints; the posterior probability is the segmentation’s share of the enumerated posterior. The leading segmentation is the single breakpoint at development year 2.

Breakpoint set	K	Posterior probability
$\{2\}$	1	0.682
$\{2, 7\}$	2	0.091
$\{2, 6\}$	2	0.063
$\{2, 5\}$	2	0.048
$\{2, 4\}$	2	0.040
$\{2, 3\}$	2	0.036

Figure 3 shows the model-averaged growth rate at each development year, with the leading segmentation overlaid as a step function. The growth rate is identified precisely at development years 1 and 2, where the in-sample cells are dense, and less precisely in the tail, where the model-averaged interval widens.




Figure 3: Per-year development growth rate, the change in the development log-effect from one year to the next. Purple markers are the model-averaged posterior mean with 90 percent credible intervals; the orange step line is the leading segmentation $B = \{2\}$, with constant rates within segments and a dotted vertical line at its breakpoint. The horizontal axis is the development year; the vertical axis is the log change per year.

4.3 Development-year effect and shares

Table 4 and Figures 4 and 5 give the model-averaged development-year effect and the implied shares of ultimate. The cumulative shares are 18.8 percent at development year 0, 50.5 percent at development year 1, and 69.8 percent at development year 2; the run-off is nearly 90 percent complete by development year 4 and approaches 99 percent by development year 8.

Table 4: Model-averaged development-year effect. β_j is the development log-effect (development year 0 set to zero) with its 90 percent credible interval; the incremental and cumulative shares are the percentage of ultimate paid in and through each development year, from the fitted 1988 cohort.

Development year	β_j (mean)	β_j (5th)	β_j (95th)	Incremental (%)	Cumulative (%)
0	0.000	0.000	0.000	18.87	18.87
1	0.442	0.401	0.483	31.64	50.51
2	-0.117	-0.156	-0.083	19.38	69.89
3	-0.671	-0.723	-0.629	11.69	81.58
4	-1.219	-1.286	-1.161	7.27	88.85
5	-1.757	-1.844	-1.661	4.53	93.38
6	-2.277	-2.400	-2.077	2.83	96.21
7	-2.759	-2.952	-2.382	1.81	98.02
8	-3.241	-3.519	-2.558	1.18	99.20
9	-3.723	-4.088	-2.713	0.80	100.00




Figure 4: Model-averaged development log-effect β_j (development year 0 set to zero), the cumulative sum of the per-year growth rates. The line is the posterior mean and the band is the 90 percent credible interval, both averaged over the breakpoint posterior. The horizontal axis is the development year; the vertical axis is the development log-effect.




Figure 5: Model-averaged incremental share of ultimate paid in each development year, from the fitted 1988 cohort. The horizontal axis is the development year; the vertical axis is the incremental share in percent, annotated above each bar.

4.4 Accident-year effect

Table 5 gives the model-averaged accident-year effect as the cohort level relative to 1997, and Figure 6 its year-on-year growth rate. The levels lie between 1.000 for accident year 1997 and 2.220 for accident year 1991, the cohort level rising to a 1991 peak and declining thereafter as the calendar-year anchor carries the nominal trend and the accident-year effect carries the residual, real, cohort scale.

Table 5: Model-averaged accident-year effect as the cohort level relative to 1997, $\exp(\alpha_i - \alpha_{1997})$, with its 90 percent credible interval.

Accident year	Level (mean)	Level (5th)	Level (95th)
1988	1.481	1.374	1.592
1989	1.682	1.558	1.810
1990	2.028	1.876	2.187
1991	2.220	2.049	2.399
1992	2.188	2.016	2.367
1993	1.993	1.835	2.157
1994	1.650	1.521	1.786
1995	1.368	1.262	1.480
1996	1.087	1.004	1.175
1997	1.000	1.000	1.000



Figure 6: Model-averaged accident-year growth rate, the year-on-year change in the accident-year effect $\Delta\alpha_i = \alpha_i - \alpha_{i-1}$ in percent per year, the increment of the first-order random walk on the cohort levels. The line is the posterior mean and the band the 90 percent credible interval; the dashed line marks zero growth. The horizontal axis is the accident year; the vertical axis is the growth rate in percent per year.

4.5 Calendar-year effect

Table 6 gives the model-averaged calendar-year effect, and Figure 7 its year-on-year growth rate. The calendar-year effect is anchored to the backbone, the cumulative logarithm of nominal gross domestic product, whose average in-sample slope is 5.4891 percent per year; the curvature is the Ornstein–Uhlenbeck deviation, which for this triangle is identified, with persistence $\phi = 0.45$ and stationary standard deviation $\sqrt{V} = 0.0212$ on the log scale, so the year-on-year growth rate departs from the anchored average by an identified, persistent amount that accumulates into the calendar curvature.

Table 6: Model-averaged calendar-year effect. γ_t is the calendar log-effect with its 90 percent credible interval; $\exp(\gamma_t)$ is the effect on the multiplicative scale. The average slope is anchored to the backbone of nominal gross domestic product, of average in-sample slope 5.4891 percent per year; the curvature is the identified Ornstein–Uhlenbeck deviation.

Calendar year	γ_t (mean)	γ_t (5th, 95th)	$\exp(\gamma_t)$
1988	−0.0210	(−0.045, 0.004)	0.979
1989	0.0540	(0.032, 0.076)	1.056
1990	0.1225	(0.099, 0.146)	1.130
1991	0.1708	(0.146, 0.196)	1.186
1992	0.2431	(0.217, 0.269)	1.275
1993	0.3097	(0.282, 0.337)	1.363
1994	0.3556	(0.330, 0.382)	1.427
1995	0.3825	(0.359, 0.406)	1.466
1996	0.4134	(0.392, 0.435)	1.512
1997	0.4590	(0.433, 0.485)	1.583




Figure 7: Model-averaged calendar-year growth rate, the year-on-year change in the calendar-year effect $\Delta\gamma_t = \gamma_t - \gamma_{t-1}$ in percent per year, with the average over the in-sample years anchored to the backbone of nominal gross domestic product, average in-sample slope 5.4891 percent per year. The line is the posterior mean and the band the 90 percent credible interval; the dotted line marks the anchored average. The horizontal axis is the calendar year; the vertical axis is the growth rate in percent per year.

4.6 Calendar forecast: the anchored backbone and a reverting deviation

The future diagonals carry calendar-year effects beyond the latest observed diagonal, and the Ornstein–Uhlenbeck treatment of Section 2.5 supplies them. The central forward path is the backbone, the cumulative logarithm of nominal gross domestic product read forward from the

published projection, and the deviation reverts to zero, so the forward calendar uncertainty is the bounded Ornstein–Uhlenbeck variance rather than one that grows without limit. Table 7 reports the projected calendar-year growth over the forecast horizon, 1998 through 2006: an annual growth of nominal gross domestic product that runs from 5.4331 percent in 1998 to 4.7320 percent over the extended diagonals 2004 through 2006 and averages 4.8070 percent, below the in-sample average of 5.4891 percent, with a cumulative calendar factor relative to 1997 that reaches 1.5413 by 2006. The projection follows the World Economic Outlook path over the six future diagonals 1998 through 2003, and the remaining diagonals 2004 through 2006 are extended at the last observed annual growth rate of 4.7320 percent, the steady-state rate at which the projection terminates.

The reversion rate of the deviation is tested for identification against the calendar curvature the triangle resolves. For the present triangle it is identified (the marginal-likelihood audit is reported in Appendix C): the curvature separates a persistent excursion from independent year-to-year departures, and the deviation carries an estimated reversion rate, with persistence $\phi = 0.45$ and stationary standard deviation $\sqrt{V} = 0.0212$ on the log scale. Forward, the deviation reverts toward the projected backbone at rate ϕ , so the forward calendar variance grows from the one-step innovation toward the stationary bound V rather than without limit, and the calendar curvature carries an identified, bounded forward uncertainty that the one-parameter form would set to zero. Figure 8 plots the in-sample backbone, the projected backbone, and the bounded deviation band.

Table 7: Projected calendar-year growth of the external anchor over the forecast horizon. The forecast year is the future calendar-year diagonal. The projected growth is the annual logarithmic growth of nominal gross domestic product read from the published projection, in percent. The cumulative calendar factor is the exponential of the cumulative projected growth relative to 1997, the multiplicative factor by which the calendar-year level scales from the last observed diagonal to the forecast year. The projection follows the World Economic Outlook path over 1998 through 2003; the diagonals 2004 through 2006 are extended at the last observed annual growth rate. Source: International Monetary Fund World Economic Outlook Database (October 2013), United States nominal gross domestic product, series NGDP.

Forecast year	Projected growth (%)	Cumulative factor
1998	5.4331	1.0558
1999	6.1505	1.1228
2000	6.2562	1.1953
2001	3.2090	1.2343
2002	3.2858	1.2755
2003	4.7320	1.3373
2004	4.7320	1.4021
2005	4.7320	1.4701
2006	4.7320	1.5413




Figure 8: Calendar-year effect: the anchored backbone in sample, the projected backbone forward, and the bounded Ornstein–Uhlenbeck deviation band. The horizontal axis is the calendar year, 1988 through 2006; the vertical axis is the calendar-year effect on the log scale, the cumulative log calendar level relative to 1988. The solid line is the backbone, the cumulative log nominal gross domestic product, actual through 1997 and projected from 1998, the projection the World Economic Outlook path over the six future diagonals 1998 through 2003, extended over 2004 through 2006 at the last observed annual growth rate. The shaded band is the forward deviation, the bounded Ornstein–Uhlenbeck forecast-error standard deviation, which widens with horizon from the one-step innovation toward the stationary standard deviation $\sqrt{V} = 0.0212$ as the reverting deviation accumulates at persistence $\phi = 0.45$. Source: BLISS v0.8.0 Ornstein–Uhlenbeck fit; anchor International Monetary Fund World Economic Outlook Database (October 2013), United States nominal gross domestic product, series NGDP.

4.7 The early-1990s calendar departure: a closed breakpoint

The calendar-year panel of Figure 12 (Section 5.4) carries a departure over the early-1990s diagonals: the standardized calendar-diagonal residual averages -0.665 over the diagonals 1988 through 1992 and turns positive over 1993 and 1994 before returning toward zero, the arc the duration channel of Section 3 anticipates.²⁰ This departure is the lagged calendar-year effect of the 1990–91 recession, carried by a closed calendar-year breakpoint on the 1993 diagonal. This subsection sets out the breakpoint, the data that place it, and the robustness of the total loss reserve estimate to the choice between the breakpoint and the reverting deviation of Section 2.5 as its representation.

The reserve estimate does not require the anchor to measure the calendar-year growth rate at the correct magnitude. A difference between the true calendar-year growth rate and the backbone that is constant in time is a linear term in the cumulative calendar level; under the age–period–cohort restriction of Section 2.2 it is reabsorbed by the accident-year and development-year effects, in

²⁰Lipton, Robertson, and Porter, “Temporary Total Disability Indemnity Benefit Duration,” Figure 2 of the source.

sample and in the projection, so the loss reserve estimate is unchanged. This is not given a priori: the growth rate of nominal gross domestic product need not measure the sum of medical-price inflation and utilization growth at the correct magnitude. What the backbone is required to carry is the departure of the calendar-year growth rate from a constant rate, not the rate itself. The early-1990s departure is one such, and the growth rate of nominal gross domestic product does not capture it: the workers'-compensation response runs through claim duration and lags the cycle. The median duration of unemployment continued to rise after the 1990–91 recession ended and held an elevated level through 1994 (Figure 9), and the calendar-year departure follows that lag onto the 1993 and 1994 diagonals rather than onto the recession.




Figure 9: Median duration of unemployment in the United States, in weeks, annual averages of the monthly seasonally adjusted series, 1988 through 2000. The horizontal axis is the calendar year; the vertical axis is the annual mean of the median number of weeks unemployed. The blue line connects the annual values; the gray band marks the National Bureau of Economic Research recession of July 1990 to March 1991. The median duration rose after the recession ended and held an elevated level through 1994, its highest annual value of the episode, before declining. Source: U.S. Bureau of Labor Statistics, Median Weeks Unemployed (series UEMPMED), retrieved from the Federal Reserve Bank of St. Louis (FRED); recession dating, National Bureau of Economic Research.

A calendar-year breakpoint is a change in the slope of the calendar-year effect at a candidate diagonal, a second-difference feature and so identified under the age–period–cohort restriction of Section 2.2. The candidate diagonals are 1990 through 1994, set by the three-increment criterion of Section 2.4 carried to the calendar axis: the hinge changes slope between a candidate year and the next, so a break at year t is carried by the increments $\Delta_{\gamma_{t+1}}$ through Δ_{γ_9} , and three or more requires t at or before 1994. The three post-break increments then fall on the cell-rich late diagonals 1995, 1996, and 1997, which carry eight, nine, and ten cells. A departure later than 1994 has too

few post-break diagonals to separate from a transient excursion and is carried by the reverting deviation rather than by a break. The breakpoint is carried on a hinge orthogonalized to the constant, the linear trend, and the books-close diagonal, so the deviation it describes rises at the break and returns to the anchored trend by the last observed diagonal—a closed arc rather than a permanent shift. The closed arc holds the latest accident-year level free of the break and confines the forward calendar to the backbone. The breakpoint occupies one direction of the calendar curvature space and the reverting deviation carries the orthogonal complement, so the two are estimated in one model without competing for the same signal.

Table 8 reports the posterior probability of a breakpoint and its location across the magnitude-prior scale σ_C . Three readings hold across the range. The location is robust: the 1993 diagonal is the most probable break position at every prior scale. The in-sample fit improves modestly: the base scale falls from $\sigma = 0.0250$ under the reverting deviation alone to between 0.0248 and 0.0222 as the breakpoint prior widens, and the calendar-diagonal residual flattens toward zero (Figure 13). The posterior weight on the break rises with the prior scale, from 0.49 at a tight prior to 0.98 at a loose one: ten diagonals fix the location and the fit but not the magnitude, so the weight the data place on the break is set in part by the prior. The economic rationale of Section 3, the lagged calendar-year effect of the 1990–91 recession, identifies the departure as the closed breakpoint, with the reverting deviation its observationally-close alternative on the ten diagonals.

Table 8: Posterior probability of a calendar-year breakpoint and its location, by the scale of the breakpoint magnitude prior. The column σ_C is the prior standard deviation of the breakpoint magnitude on the log scale, the Gaussian scale assigned to the orthogonalized hinge coefficient. The probabilities of no break and of a break are computed by exact enumeration over the candidate diagonals 1990 through 1994, with the development segmentation held at the leading structure $\{2\}$ of Section 4.2 to isolate the sensitivity to σ_C ; at $\sigma_C = 0.050$ the joint enumeration over all 64 development segmentations returns the same reading, a break probability of 0.80 placed at 1993. The reverting deviation of Section 2.5 carries the orthogonal complement of the calendar curvature in every configuration. The most probable location is the candidate diagonal carrying the largest posterior weight conditional on a break. The base scale is the in-sample observation scale σ at the leading segmentation. Across the range the reserve estimate at the leading segmentation stays within half a percent of the reverting-deviation figure of Section 6, and the joint enumeration over the development segmentation places the model-averaged reserve estimate at 0.33 percent below it. Source: BLISS v0.8.0 calendar-breakpoint enumeration.

σ_C	P(no break)	P(break)	Most probable location	Base scale σ
0.030	0.51	0.49	1993	0.0248
0.050	0.20	0.80	1993	0.0231
0.075	0.06	0.94	1993	0.0224
0.100	0.03	0.97	1993	0.0222
0.150	0.02	0.98	1993	0.0222

The total loss reserve estimate is robust to the representation; the accident-year reserve estimates are not. Integrating over the development segmentation and the calendar break jointly, the model-averaged reserve estimate is USD 312,199.3 thousand against the USD 313,239.6 thousand of Section 6, a difference of 0.33 percent, and the latest-accident-year ultimate moves by the same 0.33 percent. Beneath that total the closed breakpoint draws development weight onto the leading, shorter-run-off segmentation and reallocates across the cohorts: the tail-weighted reserve estimates of the 1989 through 1992 accident years fall by 3.2 to 4.4 percent while the 1995 and 1996 estimates rise by about 1 percent, the cohort movements netting to the 0.33 percent

total (Table 9). The closed breakpoint returns to the backbone by the last observed diagonal and confines the forward calendar to the backbone, so it does not enter the projection. The reserve estimate of Section 6 is carried under the reverting deviation of Section 2.5; representing the early-1990s departure by the closed breakpoint instead moves the total by the 0.33 percent above, the two the ten diagonals do not separate.

Table 9: Loss reserve estimate by accident year under the two representations of the early-1990s calendar-year departure. The reserve estimate is the analytic mean of the outstanding loss, the sum of the future incremental means each carrying the forecast calendar-year effect of its payment year. Under the reverting deviation the calendar curvature is carried by the Ornstein–Uhlenbeck deviation of Section 2.5 alone, and the column reproduces the per-accident-year outstanding loss of Table 13; under the model-averaged representation the estimate integrates over the development segmentation and the closed calendar-year breakpoint jointly. The change is the model-averaged estimate less the reverting-deviation estimate, in percent of the reverting-deviation estimate. The 1988 accident year is fully developed and carries no outstanding loss. Amounts are in USD thousands; the reverting-deviation total is the analytic mean of Section 6, USD 313,239.6 thousand, and the model-averaged total is USD 312,199.3 thousand, the displayed one-decimal entries summing to within a rounding unit of each. Source: BLISS v0.8.0.

Accident year	Reserve estimate (USD thousands)		Change (percent)
	Reverting deviation	Model-averaged	
1988	0.0	0.0	—
1989	1,260.1	1,204.2	−4.4
1990	3,923.6	3,757.8	−4.2
1991	8,514.0	8,185.9	−3.9
1992	15,083.3	14,594.1	−3.2
1993	23,964.5	23,438.6	−2.2
1994	34,370.3	34,140.7	−0.7
1995	49,296.3	49,703.7	+0.8
1996	67,940.1	68,724.4	+1.2
1997	108,887.3	108,450.0	−0.4
Total	313,239.6	312,199.3	−0.33

4.8 The value of the calendar-year anchor

The calendar-year axis contributes to the in-sample fit through its curvature; the anchor that sets the calendar slope contributes through the decomposition it provides, not through the fit. Constraining the calendar-year effect to zero—removing the axis, so the accident-year and development-year effects carry the diagonal movement alone—raises the in-sample base scale from 0.0247 to 0.0281 and reinstates the early-1990s residual arc that the calendar curvature absorbs: the standardized calendar-diagonal residual averaged over the 1988 through 1992 diagonals moves to −0.788, from −0.665 under the reverting deviation of Section 2.5. The reverting deviation carries that improvement. The slope is a separate matter. Refitting with a flat anchor—the backbone slope set to zero, so the accident-year and development-year effects absorb the diagonal trend while the reverting deviation is retained—gives a central reserve of USD 313,577.1 thousand against the anchored USD 313,239.6 thousand, a difference of 0.1 percent, at an in-sample base scale of 0.0247, identical to the anchored fit. The fit and the reserve estimate are little changed, the slope re-absorbed by the accident-year and development-year effects in sample and in projection. What the slope changes is the separation of run-off from inflation: under the flat anchor the development-year effect absorbs the diagonal trend, its development-year-9 log-effect reaching

−3.226 against the anchored −3.723, a gap of 0.497 log units that is the inflation the flat development effect must carry; and the anchor is the vehicle that makes the forward-inflation uncertainty of Section 4.9 quantifiable. The curvature earns the fit; the anchor’s value is identification, not point accuracy.

The base-scale and reserve comparisons of the preceding paragraph are descriptive: they report the direction and the magnitude of the in-sample change as the calendar-year axis is removed or its slope is flattened, and the base scale declines as parameters are added whether or not the added structure is warranted. The marginal likelihood supplies the inferential counterpart, penalizing the added parameters. For each of the three calendar specifications the marginal evidence is the Laplace integral over the mean-effect parameters—the reverting deviation’s curvature penalized through the Laplace determinant and the Ornstein–Uhlenbeck prior of Section 2.5—summed over the development segmentation under the breakpoint-count and location priors, by the same construction the segmentation posterior uses.²¹ The observation noise scale enters each specification through the empirical-Bayes loop, which optimizes it per specification rather than integrating it; holding the noise shape common across the three specifications removes that scale from the comparison, so the Gaussian normalizer cancels and the evidence ratio isolates the calendar mean structure, with the base scale profiled within each specification. The ratio is reported at two reference noise shapes, the full fit’s and the no-calendar fit’s, so the comparison does not depend on which specification’s noise calibration is adopted. Table 10 reports the three natural-log Bayes factors and their grades.

Table 10: Marginal-likelihood comparison of the three calendar specifications, as natural-log Bayes factors. Each specification’s marginal evidence is the Laplace integral over its mean-effect parameters—the reverting deviation’s curvature penalized through the Laplace determinant and the Ornstein–Uhlenbeck prior—summed over the 64 development segmentations under the breakpoint-count and location priors. The observation noise shape is held common across the three specifications, so the Gaussian normalizer cancels and the ratio isolates the calendar mean structure; the base scale is profiled within each specification. The log Bayes factor is reported at two reference noise shapes, the full fit’s and the no-calendar fit’s. The Bayes factor is the exponential of the full-fit-shape log Bayes factor, for the first specification against the second; a value below one favors the second. The evidence grade follows Kass and Raftery, reading twice the natural-log Bayes factor: 0 to 2, not worth more than a bare mention; 2 to 6, positive; 6 to 10, strong; above 10, very strong. The curvature comparison is the flat anchor against no calendar-year effect (the reverting deviation alone); the slope comparison is the full specification against the flat anchor (the anchored backbone slope alone); the axis comparison is the full specification against no calendar-year effect (the whole calendar-year axis).

Comparison	Log Bayes factor		Bayes factor	Evidence grade
	Full shape	Null shape		
Curvature	+1.110	+1.113	3.03	Positive
Slope	−0.364	−0.364	0.69	Bare mention
Calendar-year axis	+0.746	+0.749	2.11	Bare mention

The curvature comparison yields positive evidence: the reverting deviation is favored once its parameters are penalized, with a Bayes factor of 3.0. The slope comparison yields negative evidence—the flat anchor is favored—the slope, aliased with the accident-year and development-year effects under the age–period–cohort identity, contributing a parameter cost without a fit gain; this is the

²¹Kass and Raftery, “Bayes Factors,” 777.

evidence counterpart of the 0.1 percent reserve difference and the identical base scale of the preceding paragraph. The whole-axis evidence is not worth more than a bare mention, the curvature's positive evidence partly offset by the slope's cost, and the flat anchor is the single best-supported specification. The two readings are consistent: the descriptive base-scale rise of 13.8 percent measures the in-sample improvement attributable to the curvature, and the Bayes factor of 3.0 measures that improvement net of the curvature's parameter cost. The evidence locates the calendar-year effect's contribution in its curvature, and corroborates the anchor slope's value as the identification of Section 2.3, not as a contribution to fit.

4.9 Forward calendar uncertainty

The forward calendar variance combines two independent sources: the Ornstein–Uhlenbeck deviation, which reverts toward the projected backbone, and the forecast error of the anchor itself. At horizon h the calendar forecast-error variance is

$$V_{\text{cal}}(h) = V(1 - \phi^{2h}) + \sigma_a(h)^2, \quad (9)$$

the first term the h -step Ornstein–Uhlenbeck forecast-error variance conditional on the last observed deviation (Section 2.5), which grows from the one-step innovation toward the stationary variance V , and the second the variance of the anchor's cumulative-growth forecast error at horizon h . The anchor forecast-error standard deviation $\sigma_a(h)$ is estimated from the World Economic Outlook October-vintage cumulative-growth forecast errors over 2006 through 2025, on a growth-error basis that removes the level bias of data revisions. Table 11 reports the decomposition. The Ornstein–Uhlenbeck term, identified at $\sqrt{V} = 0.0212$, carries most of the one-step uncertainty—85 percent at the first horizon—and the anchor forecast error carries the rest, the anchor term growing with horizon to carry the majority by the second year.

Table 11: Forward calendar forecast-error standard deviation by horizon, decomposed into the Ornstein–Uhlenbeck deviation and the anchor forecast error. The horizon h is the number of future diagonals. The deviation term is $\sqrt{V(1 - \phi^{2h})}$, the h -step Ornstein–Uhlenbeck forecast-error standard deviation at persistence $\phi = 0.45$ and stationary standard deviation $\sqrt{V} = 0.0212$; the anchor term $\sigma_a(h)$ is the standard deviation of the World Economic Outlook cumulative-growth forecast error at horizon h , October vintages 2006 through 2025; the combined term is $\sqrt{V(1 - \phi^{2h}) + \sigma_a(h)^2}$. The last column is the Ornstein–Uhlenbeck share of the combined variance, in percent. All standard deviations are on the log scale.

Horizon h	OU deviation	Anchor $\sigma_a(h)$	Combined	OU share (%)
1	0.0189	0.0080	0.0205	84.9
2	0.0207	0.0280	0.0349	35.3
3	0.0211	0.0421	0.0471	20.0
4	0.0211	0.0540	0.0579	13.3
5	0.0212	0.0669	0.0702	9.1
6	0.0212	0.0678	0.0710	8.9

5 Diagnostics and Validation

This section reports the fit of the model to the observed triangle, the residual structure across the accident, development, and calendar directions, and the agreement between the reversible-jump sampler and the exact enumeration.

5.1 Goodness of fit

Table 12 reports the fit statistics. The observation model is heteroskedastic (Section 7.1): the cell standard deviation on the log-cumulative scale is the base scale times the cumulation-aware multiplier, $\sigma\sqrt{v_{ij}}$, and ranges from 0.0412 at development year 0 to 0.0184 at development year 9 (Figure 16). The model-averaged base scale is $\sigma = 0.0247$, and the leading segmentation fits at $\sigma = 0.0250$. The standardized cumulative residuals, each the cumulative log-deviation divided by its own cell standard deviation, have a mean near zero, a standard deviation of 0.992 near the unit value the variance model implies, and a maximum absolute value of 2.65 standard deviations.

Table 12: Breakpoint-model fit statistics. The base observation scale σ is on the log-cumulative scale, model-averaged and for the leading segmentation; the segment-rate count is the posterior-mean number of development growth-rate regimes less one (the posterior-mean number of breakpoints). The cell standard deviation is $\sigma\sqrt{v_{ij}}$, the base scale times the cumulation-aware multiplier of Section 7.1; the standardized residual is the cumulative log-deviation divided by its cell standard deviation $\sigma\sqrt{v_{ij}}$.

Statistic	Value
Observed cells	55
Base observation scale σ , model-averaged	0.0247
Base observation scale σ , leading segmentation	0.0250
Segment rates (posterior mean)	2.36
Cell SD $\sigma\sqrt{v_{ij}}$, development-year range	0.0184–0.0412
Standardized residual minimum	−2.65
Standardized residual maximum	2.55
Standardized residual mean	0.006
Standardized residual standard deviation	0.992

5.2 Fitted versus observed cumulative

Figure 10 compares the model-averaged fitted cumulative with the observed cumulative by accident year. The fitted cumulative remains close to the observed cumulative across the development range on the logarithmic scale.




Figure 10: Model-averaged fitted versus observed cumulative paid losses by accident year, one panel per accident year. The horizontal axis is development year (0 through 9); the vertical axis is the cumulative paid loss in USD thousands on a logarithmic scale, shared across panels. The blue line is the model-averaged fitted cumulative mean; orange markers are observed cumulative payments. The bottom row holds the 1996 and 1997 panels in its first two cells; the third cell carries the legend and reading guide, and the fourth is empty.

5.3 Log-deviation triangle

Figure 11 displays the cell-level cumulative log-deviations, $\log C_{ij} - \log M_{ij}$. The deviations are small relative to the fitted cumulative values; the directional residual smoothers of Section 5.4 resolve a mild early-1990s calendar-year departure that the cell-level triangle does not make apparent.




Figure 11: Triangle of cumulative log-deviations, $\log C_{ij} - \log M_{ij}$. Rows are accident years; columns are development years. Cell color follows a diverging scale from blue (observed below fitted) through white (zero) to red (observed above fitted), centered at zero, with the value printed in each cell. M_{ij} is the model-averaged fitted cumulative mean. Blank cells lie below the latest observed diagonal.

5.4 Residuals by accident, development, and calendar direction

Figure 12 reports the standardized cumulative residuals against each of the three directions of the triangle: accident year, development year, and calendar year. In the accident-year and development-year panels the smoother lies near zero with no sustained departure, so the accident-year levels and the development-year regimes each absorb the variation along their own direction. The calendar-year panel carries a departure over the early-1990s diagonals: the smoother runs from below the zero line at 1988 up to a positive arc over 1993 and 1994 and back toward zero by 1997, the feature that the duration channel of Section 3 describes. Section 4.7 identifies this departure as a closed calendar-year breakpoint, the lagged effect of the 1990–91 recession, and reports that the total loss reserve estimate does not depend on whether the breakpoint or the reverting deviation of Section 2.5 carries it. Figure 13 sets the calendar-diagonal residual under the reverting deviation against the residual under a closed-bump break at the 1993 diagonal: the break absorbs the step and the arc flattens toward zero. Because the three directions are linearly dependent, $t = i + j$, unmodeled structure in one direction would surface in the residuals along that direction.




Figure 12: Standardized cumulative residuals by accident-year, development-year, and calendar-year direction. Each panel plots the same 55 cell residuals against a different direction: accident year (left), development year (center), and calendar year (right). The vertical axis, shared across panels, is the standardized residual, the cumulative log-deviation (log observed minus log model-averaged fitted cumulative) divided by the model-averaged observation standard deviation. Blue markers are the residuals; the red line in each panel is a locally weighted scatterplot smoothing line (LOWESS; fitted without robustness reweighting) with span selected per panel by 5-fold cross-validation under the one-standard-error rule (selected span 0.95 in each panel); the dashed line marks zero.




Figure 13: Standardized calendar-diagonal residuals under the reverting deviation alone (left) and with a closed-bump calendar break at the 1993 diagonal (right), magnitude-prior scale $\sigma_C = 0.15$. Each panel plots the same 55 cell residuals against the calendar year; the vertical axis, shared across panels, is the standardized residual, the cumulative log-deviation (log observed minus log model-averaged fitted cumulative) divided by the cell standard deviation. Blue markers are the cell residuals; the red line is a locally weighted scatterplot smoothing line (LOWESS; fitted without robustness reweighting) with span selected per panel by 5-fold cross-validation under the one-standard-error rule (selected span 0.95 in each panel); the dashed line marks zero. The left panel carries the early-1990s arc of Figure 12; under the break the smoother lies near zero across the diagonals, the closed bump absorbing the early-1990s step while the reverting deviation carries the orthogonal complement. Source: BLISS v0.8.0 calendar-breakpoint fit.

5.5 Collapsed reversible jump against exact enumeration

The collapsed chain uses four runs over the segmentation space from over-dispersed starts. It reproduces the exact enumeration: the posterior over the number of breakpoints (Table 1) and the certainty of a breakpoint at development year 2 (Table 2) match, and the placement of a second breakpoint, where one is present, mixes freely across the later development years. The chain collapses the continuous parameters analytically and moves only over the discrete segmentation, which removes the slow trans-dimensional component that a joint sampler carrying the continuous parameters inside the dimension-changing move would incur. The recovered marginals agree with the enumeration to within half a unit in the third decimal place—the posterior over the number of breakpoints, the development-year-2 certainty, and the leading segmentations $\{2\}$ at 0.682 and $\{2, 7\}$ at 0.091. The outstanding loss and the ultimate are Rao-Blackwellized over the enumerated segmentations, so they carry the exact segmentation weighting rather than chain visit frequencies. The exact enumeration and the collapsed chain agree on the reported structure, and the enumeration backstops the chain where the candidate set is small enough to enumerate.

5.6 Sensitivity to the rate prior

The development-rate scale σ_r enters the comparison of segmentations through the parsimony penalty in the marginal likelihood. The hyperprior is centered at the data-equivariant scale $\sigma_r = 0.5782$, the root-mean-square of the saturated free-rate development slopes taken over the development years that at least four accident years observe (Section 2.7). The marginal likelihood of the development model is flat in σ_r over a range around this center, which is the reason the rate variance is integrated against its hyperprior rather than fixed at a point: the integration averages over the flat region instead of conditioning on one value within it. Over that range the development-year-2 breakpoint is certain and the single-breakpoint structure is unchanged, so the number and the location of the breakpoints are governed by the data and not by the rate scale.

6 Loss Reserve Estimate and Predictive Distribution

This section reports the completed projection by accident year, the posterior distribution of the total outstanding loss, the dependence of the loss reserve estimate on the assumed calendar-year trend, and the predictive distribution of the 1997 accident-year ultimate with its decomposition by source.

6.1 Completed projection

Table 13 reports the projection by accident year with the Ornstein–Uhlenbeck calendar forecast carried into the future diagonals. The outstanding loss for each accident year is the analytic mean of the sum of its future incremental means, each carrying the calendar-year effect of its payment year. The estimated total outstanding loss is USD 313,239.6 thousand—the analytic mean of the outstanding-loss distribution, which is additive across accident years—on a projected ultimate of USD 1,748,029.6 thousand.

Table 13: Ultimate projection by accident year with the Ornstein–Uhlenbeck calendar forecast. Latest paid is the observed cumulative payment along the 1997 diagonal; the outstanding loss is the analytic mean of the sum of future incremental means $m_{ij} = \exp(\mu_0 + \alpha_i + \beta_j + \gamma_{i+j})$, each carrying the forecast calendar-year effect of its payment year; projected ultimate is latest paid plus outstanding loss. The per-accident-year means are additive and sum to the total outstanding-loss analytic mean at full precision; the displayed one-decimal entries can sum to within a rounding unit of the displayed total. Amounts are in USD thousands.

Accident year	Latest paid	Outstanding loss	Projected ultimate
1988	125,049	0.0	125,049.0
1989	147,358	1,260.1	148,618.1
1990	187,760	3,923.6	191,683.6
1991	213,396	8,514.0	221,910.0
1992	213,904	15,083.3	228,987.3
1993	193,676	23,964.5	217,640.5
1994	151,081	34,370.3	185,451.3
1995	111,268	49,296.3	160,564.3
1996	66,033	67,940.1	133,973.1
1997	25,265	108,887.3	134,152.3
Total	1,434,790	313,239.6	1,748,029.6

6.2 Posterior of the total outstanding loss

The total outstanding loss is the sum of the projected future incremental amounts,

$$R = \sum_{i+j>9} \exp(\mu_0 + \alpha_i + \beta_j + \gamma_{i+j}), \quad (10)$$

with the future calendar-year effect taken from the Ornstein–Uhlenbeck forecast. Each posterior draw of the breakpoint structure, the segment rates, the accident-year levels, and the calendar curvature, together with a draw of the calendar forecast, yields one draw of the outstanding loss. The loss reserve estimate—the analytic mean of the total outstanding loss—is USD 313,239.6 thousand, the closed-form posterior mean of Section 6.4, additive across accident years. The posterior median is USD 300,960.4 thousand, with a 90 percent credible interval of USD 278,641.5 thousand to USD 339,440.1 thousand and a geometric standard deviation of 1.0618 (log-scale standard deviation 0.0600); the interval is the misspecification-robust posterior of Section 7.2, which carries the cumulation correlation that the diagonal observation model sets aside, while the central estimate is held on the penalized posterior. Under the data-equivariant development-rate prior the mean and the median are within about 4 percent; the mean is the reported central estimate, and the median, the interval, and the geometric standard deviation describe the distribution. Figure 14 shows the posterior density.




Figure 14: Posterior density of the total outstanding loss, integrating over the breakpoint structure, the segment rates, the accident-year levels, the calendar-year curvature, and the Ornstein–Uhlenbeck calendar forecast. The density is a kernel estimate of 30,000 unweighted draws at the posterior process variance; the solid red line marks the analytic mean, the dashed line the median, and dotted lines the 5th and 95th percentiles; the horizontal axis is truncated at the bounds shown because the distribution is right-skewed. The horizontal axis is the total outstanding loss in USD thousands; the vertical axis is posterior density.

6.3 Sensitivity of the loss reserve estimate to a calendar-year shock

The calendar-year effect is anchored to the level of nominal gross domestic product (Section 2.3), and the forward calendar path is the projected backbone (Section 2.5). The sensitivity of interest is to a departure of the realized calendar-year path from that projection. Table 14 records the outstanding loss under a calendar-year shock to the first future diagonal, calendar year 1998, of -2 , -1 , $+1$, and $+2$ percentage points on the annualized calendar-year growth rate. The shock occurs in the first quarter of the first future diagonal and dissipates geometrically with a half-life of four quarters: the growth-rate increment halves every four quarters and accumulates into the calendar-year level, so the level deviation it leaves climbs from 1.07 percentage points on the first future diagonal to a bounded plateau near 3.13 percentage points across the later diagonals. The baseline outstanding loss of USD 313,239.6 thousand (analytic mean) is the unshocked projected path. The entries are analytic means of the outstanding-loss distribution under each shocked path. The calendar quarter of the shock, its magnitude in percentage points, and the half-life in quarters are user-set parameters of the projection.

Across the four shocks the mean outstanding loss runs from USD 304,792.4 thousand at -2 percentage points to USD 321,934.0 thousand at $+2$ percentage points, a move of about 2.8 percent of the outstanding loss for each two-percentage-point shock. The growth-rate increment decays within about two years, but the level deviation it leaves applies to every subsequent forecast diagonal; because the workers' compensation run-off is concentrated on the near forecast diagonals rather than spread across the horizon, the reserve response is that accumulated level shift applied mainly to the near outstanding loss. The upward asymmetry between the $+2$ and -2 entries—a rise of 8,694.4 against a fall of 8,447.2 from the baseline—reflects the convexity of the exponential.

Table 14: Sensitivity of the loss reserve estimate to a calendar-year shock on the first future diagonal. The shock is applied to the annualized calendar-year growth rate at the first quarter of calendar year 1998 and dissipates geometrically with a half-life of four quarters, the growth-rate increment accumulating into the calendar-year level. The baseline shock of zero is the unshocked projected path. Entries are analytic means of the outstanding-loss distribution under each shock. The total outstanding loss and the total ultimate are in USD thousands; total ultimate is paid-to-date of 1,434,790 plus the outstanding loss.

Shock (percentage points)	Total outstanding loss	Total ultimate
-2	304,792.4	1,739,582.4
-1	308,985.5	1,743,775.5
0	313,239.6	1,748,029.6
$+1$	317,555.4	1,752,345.4
$+2$	321,934.0	1,756,724.0

6.4 The 1997 ultimate and the structural component

The predictive distribution of the 1997 accident-year ultimate is the observed development-year-0 amount plus the projected run-off, carried through observation, parameter, structural, and calendar-forecast uncertainty. Its analytic mean is USD 131,235.0 thousand, the observed development-year-0 amount of 25,265 plus the 1997 outstanding loss of USD 108,887.3 thousand from Table 13, and its median is USD 131,054.0 thousand, with a 90 percent predictive interval of USD 121,934.3 thousand to USD 141,221.9 thousand, a geometric standard deviation of 1.0456 (log-scale standard deviation 0.0446), and a 99th percentile of USD 145,838.5 thousand. The predictive distribution of the 1997 ultimate is comparatively tight: development year 0 is 18.9 per-

cent of the ultimate for this triangle, so a larger share is observed than for a slow line and the projection carries less of the total. It is reported on the same mean, median, and log-scale basis as the outstanding loss. Figure 15 shows the predictive density accumulating by source, from observation noise alone, through the addition of parameter and structure, to the full density that adds the calendar forecast.

The predictive variance separates into a within-model part and a structural part by the law of total variance,

$$\text{Var}(\log U) = \underbrace{\mathbb{E}_B [\text{Var}(\log U \mid B)]}_{\text{within-model}} + \underbrace{\text{Var}_B (\mathbb{E}[\log U \mid B])}_{\text{structural}}, \quad (11)$$

where B is the segmentation. On the log scale the dispersion separates by source without dependence on a second moment on the original, linear scale. The full predictive log-scale standard deviation is 0.0446. The point-forecast component, before observation noise—parameter, structure, and calendar forecast—is 0.0437, and the observation component is 0.0089; the point-forecast component carries nearly the whole dispersion. The observation component is small because the observation scale of the 1997 ultimate is the cumulation-aware accumulation of the incremental scale of Section 7.1, each increment entering at the square of its share of the accumulated cumulative amount, so the projected run-off contributes observation variance at a fraction of the incremental scale. The two components combine in quadrature to the full dispersion up to a small positive covariance, the accumulated observation scale depending on the same draw as the point forecast. The point-forecast component is the larger: the 1997 ultimate is almost entirely a projection and carries the projection’s parameter, structure, and calendar uncertainty. Within the point forecast the calendar-forecast contribution is small, the identified but bounded reverting calendar deviation carrying little forward uncertainty over the 1997 cohort’s short remaining run-off, so parameter and structure account for nearly all of it. The structural term of equation (11)—the part attributable to uncertainty about the run-off structure itself—is a small share of the total: the leading segmentations imply similar run-off, so the within-model component carries most of the point-forecast dispersion. Table 15 reports the predictive summary and the log-scale dispersion by source.

Table 15: Predictive distribution of the 1997 accident-year ultimate, and the predictive standard deviation by source. Quantiles and the median are in USD thousands. The log-scale standard deviation is the standard deviation of the natural logarithm of the predictive draws and the geometric standard deviation is its exponential. The two log-scale source rows are the observation-noise component and the point-forecast component (parameter, structure, and calendar forecast, before observation noise); the point-forecast component dominates the ultimate, whose run-off is projected. Computed from 30,000 unweighted draws at the posterior process variance.

Quantity	Value
Median	131,054.0
5th percentile	121,934.3
25th percentile	127,190.0
75th percentile	135,098.6
95th percentile	141,221.9
99th percentile	145,838.5
Log-scale standard deviation	0.0446
Geometric standard deviation	1.0456
Log-SD, observation only	0.0089
Log-SD, point forecast (parameter, structure, calendar)	0.0437




Figure 15: Predictive density of the 1997 accident-year ultimate, accumulated by variance source, as kernel estimates of unweighted draws at the posterior process variance over the central region. The blue dashed curve carries observation (process) variability only at the central point forecast; the green dash-dot curve adds parameter and breakpoint-structure uncertainty, with the calendar held at its forecast mean; the orange solid curve adds calendar-forecast uncertainty and is the full predictive density. The parameter-and-structure curve and the full curve nearly coincide, the calendar-forecast contribution being small. The solid red line marks the analytic mean; dotted vertical lines mark the 5th percentile, median, and 95th percentile of the full distribution. The horizontal axis is the ultimate loss in USD thousands; the vertical axis is predictive density.

The central estimate is the analytic posterior mean of the outstanding loss, computed in closed form. Conditional on a segmentation each future cell is log-normal on the log scale, so its posterior mean is the exponential of the fitted linear predictor plus one-half its posterior variance; the cell means sum to the outstanding-loss mean and are averaged over the enumerated segmentations by their posterior weights. The construction is deterministic and additive across accident years, and it does not depend on the stability of a sample mean of exponentials. The dispersion is reported on the log scale, by predictive quantiles, and by the geometric standard deviation rather than by an original-scale standard deviation. The predictive variance on the original scale is not guaranteed finite: the observation-noise multiplier's second moment is the expectation of $\exp(\sigma^2)$ under the scaled inverse- χ^2 posterior of σ^2 , which the inverse-gamma right tail does not damp, so that integral diverges, while the log-scale dispersion is finite throughout.

The data-equivariant development-rate prior holds the closed-form mean close to the median. Under a diffuse prior a weakly identified development-year rate—a late development year informed by few cells, or a candidate breakpoint the data do not constrain—reverts toward a wide prior on the log scale, and the half-variance term of the log-normal mean inflates the corresponding cell; matching the prior scale to the observed development-rate scale (Section 2.7) removes that inflation. The mean then sits within 2 percent of the median for the breakpoint model (Section 6.2). The original-

scale standard deviation remains tail-sensitive through the observation-variance multiplier alone and is not reported; the mean, the median, the percentile interval, and the log-scale dispersion are reported throughout and exist under the proper priors carried here.²² Appendix B reports the quantiles and the log-scale dispersion.

7 Modeling Considerations

Three features of the model warrant comment: the variance of the observation noise placed on the cumulative paid amount, the cumulation correlation that the same observation model sets aside, and the inference for the most recent accident year, which is observed at a single cell.

7.1 The heteroskedastic cumulative observation model

The observation model places log-normal noise on the cumulative paid amount (equation (3)), with a variance that is not constant along the development axis. The dispersion of the cumulative residuals falls steeply with development, but that decline is in large part a cumulation artifact: as the cumulative amount grows, a fixed incremental error is a smaller fraction of it, so the cumulative residual contracts whether or not the underlying increment process changes. The observation scale is therefore estimated on the increment residual, where the artifact is absent, and carried to the cumulative cell by accumulation. The increment log-variance is a log-quadratic function of the standardized development year, $s_k^2 = \phi \exp(b d_k + c d_k^2)$, where $d_k = (k - 3.0)/2.449$ is the development year centered and scaled to its observed-cell mean and standard deviation and the profile is normalized to geometric mean one over the observed cells, so that ϕ retains its meaning as the geometric-mean observation variance. The shape is fitted by empirical Bayes—alternating the model-averaged increment residuals with the regularized profile objective and the weighted re-fit—to a fixed point at a slope $b = 0.5430$ and a curvature $c = 0.3700$. The positive curvature gives the incremental scale a shallow minimum near development year 1: the log-scale standard deviation is 0.019 at development year 0, 0.019 near development year 1, and 0.121 at development year 9, so the early increments are the most precisely determined and the late increments the least. The cumulative cell variance is the cumulation-aware accumulation of these incremental scales, $v_{ij} = \sum_{k=0}^j (m_{ik}/M_{ij})^2 s_k^2$, in which each increment enters at the square of its share of the accumulated cumulative amount, so a noisy late increment contributes to a cumulative cell only at its small share rather than at its full incremental scale. The accumulated cumulative scale falls from a log-scale standard deviation of 0.041 at development year 0 to 0.018 at the tail (Figure 16), the funnel the cumulation itself produces, lying between the artifact-laden cumulative residual above it and the rising incremental scale below. The homoskedastic model, which places a single scale on every cumulative cell, is the special case $b = c = 0$ and is reached when the increment residuals do not support a departure; on this triangle the curvature is supported, and the heteroskedastic scale is carried in the reported fit.

²²For log-normal functionals—an exponential of a linear combination of the log-scale parameters—the posterior can be proper yet lack finite moments under the variance priors in common use, including the inverse-gamma prior on the log-scale variance. The point was first noted in Arnold Zellner, “Bayesian and Non-Bayesian Analysis of the Log-Normal Distribution and Log-Normal Regression,” *Journal of the American Statistical Association* 66, no. 334 (1971): 327–330, and developed for the log-normal linear model in Enrico Fabrizi and Carlo Trivisano, “Bayesian Estimation of Log-Normal Means with Finite Quadratic Expected Loss,” *Bayesian Analysis* 7, no. 4 (2012): 975–996, and “Bayesian Conditional Mean Estimation in Log-Normal Linear Regression Models with Finite Quadratic Expected Loss,” *Scandinavian Journal of Statistics* 43, no. 4 (2016): 1064–1077, who give conditions on the variance prior under which posterior moments of a prescribed order exist. Appendix B sets out the quantile and log-scale reporting.




Figure 16: Fitted observation log-scale by development year for the State Farm Workers' Compensation triangle. The horizontal axis is the development year; the vertical axis is the standard deviation of the log observation noise. The incremental scale (orange, solid, circles) is the log-quadratic increment standard deviation $s_k = \sqrt{\phi \exp(b d_k + c d_k^2)}$, with slope $b = 0.5430$ and curvature $c = 0.3700$ fitted by empirical Bayes on the model-averaged increment residuals; it runs from 0.019 at development year 0 through a shallow minimum near development year 1 to 0.121 at development year 9. The accumulated cumulative scale (blue, dashed, squares) is the cumulation-aware accumulation $\sqrt{v_{ij}}$, $v_{ij} = \sum_{k \leq j} (m_{ik}/M_{ij})^2 s_k^2$, each increment entering at the square of its share of the accumulated cumulative amount; it is the standard deviation carried into the diagonal weighting of the in-sample fit and falls from 0.041 to 0.018. The two profiles cross near development year 3, the incremental scale lying below the accumulated scale early and above it late. Source: BLISS v0.8.0 heteroskedastic fit, obs_var dev_quadratic.

7.2 A misspecification-robust posterior for the cumulation correlation

The heteroskedastic observation model weights each cumulative cell by its accumulated variance, but within an accident year it treats the cells as independent. A row of cumulative amounts is a sequence of partial sums of a common set of increments, so its cells are positively correlated, and the diagonal weighting is a working model that sets that correlation aside. Carrying the correlation inside the likelihood would enter the row covariance as a generalized-least-squares weight and so invert it; on this triangle the cumulative cells are near-collinear, the row covariance is numerically singular, and the generalized-least-squares fit does not converge. The correlation is there-

fore carried after estimation, as a curvature adjustment to the posterior covariance of the kind a misspecification-robust, or sandwich, posterior provides.²³

The parameter covariance of each segmentation, the inverse penalized Hessian H^{-1} that the Laplace approximation supplies, is replaced by the sandwich form $H^{-1}JH^{-1}$. The meat $J = \sigma^{-2}X^TW\Sigma WX$ carries the row cumulation covariance Σ , with X the Jacobian of the fitted log cumulatives at the mode and W the diagonal accumulated-variance weighting; the covariance Σ enters only through the quadratic form $W\Sigma W$ and is never inverted, so the construction is immune to the near-collinearity that leaves the generalized-least-squares fit unestimable. The central estimate is left on the penalized posterior, so the correction adjusts the dispersion of the predictive distribution and not its location, as a robust standard error adjusts an interval without moving the point estimate it surrounds.

Table 16: Geometric standard deviation of the total outstanding loss under the naive and the misspecification-robust posterior. The naive row is the diagonal Laplace posterior; the cumulation-covariance row places the row covariance Σ in the sandwich meat and adds the cumulation correlation. Both hold the analytic central estimate of USD 313,239.6 thousand and are computed at seed 11 over 30,000 draws of the predictive distribution.

Posterior covariance	Outstanding-loss geoSD
Naive Laplace (diagonal)	1.0777
Robust, cumulation-covariance meat	1.0618

The correction tightens the loss reserve interval (Table 16): the geometric standard deviation of the total outstanding loss falls from 1.0777 under the diagonal posterior to 1.0618 once the cumulation covariance enters the meat. The loss reserve is a sum of development-rate differences, and a positive correlation among cumulative cells lowers the variance of a differenced quantity, so the independent-noise assumption over-states the reserve interval rather than under-states it; the robust posterior confirms that direction and corrects for it. The 1997 ultimate, a single cohort observed at one cell, tightens by less, from 1.0480 to 1.0456, as its row carries no within-row correlation to exploit.

The same correlation enters the evidence that weights the segmentations, where its effect runs the other way: the diagonal evidence treats the correlated cells as independent observations and so over-states the information that separates one segmentation from another, an effect a single information-ratio tempering of the evidence places near a factor of 0.29. The within-model interval is corrected here, because the sandwich covariance is a validated estimator; the between-model weights are reported as the diagonal evidence gives them, with that tempering carried as a sensitivity, because a single-factor tempering is a diagnostic and not a calibrated correction. The loss reserve estimate is in any case insensitive to the segmentation weights.

7.3 Inference for the latest accident year

The 1997 accident year is observed only at development year 0. Its level is not identified from its single cell alone; it is identified through the first-order random-walk prior of Section 2.6, which

²³Ulrich K. Müller, “Risk of Bayesian Inference in Misspecified Models, and the Sandwich Covariance Matrix,” *Econometrica* 81, no. 5 (2013): 1805–1849.

pools information from the adjacent cohorts in the manner of a shrinkage estimator.²⁴ The outstanding loss for 1997, USD 108,887.3 thousand in the mean, is the largest of any accident year, because the cohort is observed at development year 0 alone and the remaining run-off is projected, so the 1997 level carries weight in the total outstanding loss out of proportion to its single observation. The posterior level of the latest accident year therefore rests on the random-walk prior and the neighboring cohorts as much as on its own datum, and its credible interval reflects that the cohort is informed indirectly.

8 Conclusion

This paper treats the loss-development structure as an object of inference rather than a fixed specification. The development-year effect is represented as a sequence of growth-rate regimes, and the number and the location of the breakpoints are inferred by reversible jump. The age-period-cohort aliasing that would otherwise leave the calendar-year slope unidentified is resolved by anchoring the calendar-year effect to the level of nominal gross domestic product, the linear trend fixed to the anchored backbone and the curvature carried by a mean-reverting deviation and, over the early-1990s diagonals, a closed calendar-year breakpoint. The loss reserve estimate and its predictive distribution are reported with the structural uncertainty carried through, and the development-axis sampler is checked against the exact posterior obtained by enumeration.

Several limitations bound these findings. The calendar-year forecast carries a Gaussian deviation about the anchored backbone; the deviation reverts toward the projected backbone under the Ornstein–Uhlenbeck process of Section 2.5, its reversion rate tested for identification and, for this triangle, identified at persistence $\phi = 0.45$, so the forward calendar uncertainty is bounded by the stationary deviation variance and carries an estimated persistence rather than growing without limit. The observation model places log-normal noise on the cumulative amounts with a development-year-varying scale, and the cumulation dependence of the partial-sum sequence, which a diagonal likelihood sets aside, is carried into the predictive interval by a misspecification-robust posterior (Section 7.2) that tightens it, the diagonal posterior having over-stated the reserve interval. The findings rest on a single triangle.

These limitations mark the directions left to subsequent work. A skewed or heavier-tailed calendar-year deviation would accommodate a triangle whose realized calendar-year growth departs from normality; the present triangle does not reject normality, so the extension is not exercised here. Validation across additional triangles, together with a simulation study in which the breakpoint locations are known, would establish the recovery properties of the sampler beyond the single application reported here.

The paid-loss triangle is the State Farm Workers’ Compensation triangle from the NAIC Schedule P loss-reserving database compiled by Meyers and Shi, reproduced in Appendix A; the macroeconomic series is United States nominal gross domestic product from the IMF World Economic Outlook, cited in Section 3. The model, the priors, and the inference are specified in full in Section 2.

²⁴Bradley Efron and Carl Morris, “Stein’s Estimation Rule and Its Competitors—An Empirical Bayes Approach,” *Journal of the American Statistical Association* 68, no. 341 (1973): 117–130.

A Paid-Loss Triangles

Tables A.1 and A.2 reproduce the incremental and cumulative paid-loss triangles. Cells below the latest observed diagonal are unobserved and are the object of projection.

Table A.1: Incremental paid-loss triangle, in thousands of United States dollars. Rows are accident years; columns are development years. Empty cells lie below the latest observed diagonal. Source: State Farm Workers' Compensation, NAIC Schedule P loss-reserving database (Meyers and Shi).

AY	0	1	2	3	4	5	6	7	8	9
1988	22,190	38,644	24,270	15,047	8,661	6,155	3,823	2,768	1,934	1,557
1989	26,542	51,256	28,609	16,015	10,937	5,240	4,430	2,683	1,646	
1990	32,977	67,517	34,392	22,872	11,233	9,074	4,722	4,973		
1991	38,604	75,824	42,675	24,219	16,089	11,393	4,592			
1992	42,466	83,354	38,956	24,269	15,332	9,527				
1993	46,447	70,317	38,133	24,522	14,257					
1994	41,368	58,976	31,677	19,060						
1995	35,719	47,497	28,052							
1996	28,746	37,287								
1997	25,265									

Table A.2: Cumulative paid-loss triangle, in thousands of United States dollars, formed as the running sums of the incremental source integers. Rows are accident years; columns are development years. Empty cells lie below the latest observed diagonal. Source: State Farm Workers' Compensation, NAIC Schedule P loss-reserving database (Meyers and Shi).

AY	0	1	2	3	4	5	6	7	8	9
1988	22,190	60,834	85,104	100,151	108,812	114,967	118,790	121,558	123,492	125,049
1989	26,542	77,798	106,407	122,422	133,359	138,599	143,029	145,712	147,358	
1990	32,977	100,494	134,886	157,758	168,991	178,065	182,787	187,760		
1991	38,604	114,428	157,103	181,322	197,411	208,804	213,396			
1992	42,466	125,820	164,776	189,045	204,377	213,904				
1993	46,447	116,764	154,897	179,419	193,676					
1994	41,368	100,344	132,021	151,081						
1995	35,719	83,216	111,268							
1996	28,746	66,033								
1997	25,265									

B Finite-Variance Reporting of the Predictive Distribution

The closing note of Section 6.4 records that the original-scale predictive mean is finite but tail-sensitive in estimation while the original-scale predictive variance is not guaranteed finite—under the scaled inverse- χ^2 posterior of the observation variance and, more materially, under the diffuse normal prior on weakly identified development-year rates. The dispersion of the predictive distribution is therefore reported without relying on a population second moment that may not exist.

B.1 Log-scale dispersion and quantiles

The predictive distribution is summarized by the analytic mean as the central estimate and by its quantiles and log-scale dispersion, the quantiles and the log-scale dispersion existing regardless of the moment condition. Table B.1 gives the analytic mean, the original-scale quantiles of the total outstanding loss and the 1997 accident-year ultimate, the log-scale standard deviation, and the geometric standard deviation.

Table B.1: Predictive quantiles and log-scale dispersion of the total outstanding loss and the 1997 accident-year ultimate. The quantiles are in USD thousands. The log-scale standard deviation is the standard deviation of the natural logarithm of the predictive draws; the geometric standard deviation is its exponential, the multiplicative factor by which a one-standard-deviation move on the log scale scales the dollar amount. The analytic mean is the closed-form posterior mean reported as the central estimate in the body (Section 6.2 and Section 6.4); the original-scale standard deviation is omitted as tail-sensitive through the observation-variance multiplier. Quantiles and dispersion are computed from 30,000 unweighted draws at the posterior process variance, Rao-Blackwellized over the enumerated segmentations.

	Total outstanding loss	1997 ultimate
Analytic mean	313,239.6	131,235.0
1st percentile	271,197.8	118,400.4
5th percentile	278,641.5	121,934.3
10th percentile	283,063.0	123,859.1
25th percentile	290,889.9	127,190.0
50th percentile (median)	300,960.4	131,054.0
75th percentile	315,382.0	135,098.6
90th percentile	331,082.1	138,777.1
95th percentile	339,440.1	141,221.9
99th percentile	353,379.6	145,838.5
Log-scale standard deviation	0.0600	0.0446
Geometric standard deviation	1.0618	1.0456

B.2 A thin-tailed prior for the observation variance

The scaled inverse- χ^2 , equivalently inverse-gamma, posterior of σ^2 has a polynomial right tail, so the expectation of $\exp(\sigma^2)$ —the second moment of the observation-noise multiplier—diverges, and the original-scale predictive variance is not a finite population quantity. A generalized inverse Gaussian (GIG) prior on σ^2 replaces the polynomial right tail with an exponential one through a density factor $\exp(-\psi\sigma^2/2)$. Under it the moments of the original-scale predictive through order r exist when $\psi > r(r-1)$, so $\psi > 2$ secures a finite predictive variance. The inverse-gamma posterior is the $\psi \rightarrow 0$ limit, so the GIG posterior equals it multiplied by $\exp(-\psi\sigma^2/2)$. Over the range of σ^2 the draws reach, $[3.0 \times 10^{-3}, 1.6 \times 10^{-2}]$, that factor departs from one by at most 3.1×10^{-2} at $\psi = 4$, concentrated in the upper tail of σ^2 ; the thin-tailed prior mildly downweights the largest ob-

ervation variances and slightly narrows the upper predictive tail. The thin-tailed prior secures the moments of the observation-noise multiplier. The other potential source of a heavy upper tail—a diffuse prior on weakly identified development-year rates, whose exponential would carry a long right tail into the original-scale outstanding loss—is addressed directly by the calibration of the rate-prior scale to the observed development-rate scale (Section 2.7), which holds the closed-form outstanding-loss mean within about four percent of the median for the smooth model and within two percent for the breakpoint model. Under the calibrated rate prior the closed-form mean is the reported central estimate; the original-scale variance under the inverse-gamma observation-variance prior remains tail-sensitive, so the log-scale and quantile reporting of Section B.1 is the operative summary of dispersion, and the generalized inverse Gaussian prior is the route to a finite original-scale variance where one is required.

C Evidence Profile of the Deviation Parameters

The deviation parameters of Section 2.7 enter the fit as supplied constants, with persistence $\phi = 0.45$ and stationary standard deviation $\sqrt{V} = 0.0212$. This appendix audits both by marginal likelihood. For the amplitude, the full empirical-Bayes fit re-runs at each of seven half-octave amplitudes at the supplied persistence, the observation-noise shape re-equilibrating per point, and each point is read by the total marginal likelihood—the Laplace evidence summed over the 64 development segmentations under the breakpoint-count and location priors, the construction of Section 2.8—at that point’s converged shape. Each converged shape is normalized to geometric mean one over the observed cells, so the determinant term of the heteroskedastic Gaussian is zero at every point and the per-amplitude readings are comparable; the computation asserts that invariant point by point. A profile at one fixed noise shape rises monotonically in the amplitude, a larger prior amplitude buying curvature fit at that shape without the offsetting recalibration; the per-point re-equilibration supplies the trade-off. Table C.1 reports the profile.

Table C.1: Marginal-likelihood profile of the Ornstein–Uhlenbeck amplitude at persistence $\phi = 0.45$. The first column is the stationary standard deviation $\sqrt{V}$ of the deviation, on the log scale, on a seven-point half-octave grid. The second column is the total marginal log-likelihood at that amplitude relative to its maximum over the grid, in natural-log units: at each amplitude the full empirical-Bayes fit re-runs, the observation-noise shape re-equilibrates, and the marginal likelihood is the Laplace evidence summed over the 64 development segmentations under the breakpoint-count and location priors at the converged shape, which is normalized to geometric mean one over the observed cells. The third column is the converged base observation scale σ , on the log scale. Source: BLISS v0.8.0.

$\sqrt{V}$	Log marginal likelihood, relative to maximum	Converged base scale σ
0.0075	−2.642	0.0281
0.0106	−2.314	0.0276
0.0150	−1.712	0.0265
0.0212	−0.774	0.0247
0.0300	0.000	0.0222
0.0424	−1.048	0.0201
0.0600	−4.268	0.0194

The profile carries an interior maximum at $\sqrt{V} = 0.030$. The supplied amplitude sits 0.774 natural-log units below the maximum; twice the margin, 1.55, reads “not worth more than a bare mention” on the grading of Table 10.²⁵ At the supplied amplitude, the evidence for the reverting family—persistence positive, uniform over the seven-point persistence grid—against the independent-innovations boundary $\phi = 0$, the one-parameter form of Section 2.5, reads twice the natural-log Bayes factor (BF), $2 \log \text{BF} = 1.27$, by family average and 1.95 at the peak persistence $\phi = 0.60$, below the positive grade; at the profile-maximum amplitude the family reading is 2.61, clearing it.²⁶ The two readings cohere: the marginal likelihood prefers, without clearing the positive grade over the supplied value, an amplitude at which the reverting family clears the positive grade. Holding the supplied constants leaves every quantity reported in the body unchanged;

²⁵Kass and Raftery, “Bayes Factors,” 777.

²⁶The family evidence is the log-sum-exp of the per-persistence evidences minus the logarithm of the point count, a uniform prior over the grid; the persistence grid is $\phi \in \{0.15, 0.30, 0.45, 0.60, 0.75, 0.90\}$ against the boundary $\phi = 0$, read at the standing fit’s observation-noise shape, where holding the shape common cancels the determinant term by the construction of Section 4.8.

adopting the profile-maximum amplitude would rest on a distinction the evidence grades below positive. The supplied constants stand, with this appendix as their audit.

Bibliography

- Barnett, Glen, and Ben Zehnwirth. "Best Estimates for Reserves." *Proceedings of the Casualty Actuarial Society* 87 (2000): 245–321.
- Denison, D. G. T., B. K. Mallick, and A. F. M. Smith. "Automatic Bayesian Curve Fitting." *Journal of the Royal Statistical Society, Series B* 60, no. 2 (1998): 333–350.
- DiMatteo, Ilaria, Christopher R. Genovese, and Robert E. Kass. "Bayesian Curve-Fitting with Free-Knot Splines." *Biometrika* 88, no. 4 (2001): 1055–1071.
- Efron, Bradley, and Carl Morris. "Stein's Estimation Rule and Its Competitors—An Empirical Bayes Approach." *Journal of the American Statistical Association* 68, no. 341 (1973): 117–130.
- England, Peter D., and Richard J. Verrall. "Stochastic Claims Reserving in General Insurance." *British Actuarial Journal* 8, no. 3 (2002): 443–518.
- Fabrizi, Enrico, and Carlo Trivisano. "Bayesian Estimation of Log-Normal Means with Finite Quadratic Expected Loss." *Bayesian Analysis* 7, no. 4 (2012): 975–996.
- Fabrizi, Enrico, and Carlo Trivisano. "Bayesian Conditional Mean Estimation in Log-Normal Linear Regression Models with Finite Quadratic Expected Loss." *Scandinavian Journal of Statistics* 43, no. 4 (2016): 1064–1077.
- International Monetary Fund. *World Economic Outlook Database, October 2013*. Washington, DC: International Monetary Fund, 2013. United States gross domestic product at current prices, series NGDP.
- Geisser, Seymour, and William F. Eddy. "A Predictive Approach to Model Selection." *Journal of the American Statistical Association* 74, no. 365 (1979): 153–160.
- Gelman, Andrew, Xiao-Li Meng, and Hal Stern. "Posterior Predictive Assessment of Model Fitness via Realized Discrepancies." *Statistica Sinica* 6, no. 4 (1996): 733–760.
- Green, Peter J. "Reversible Jump Markov chain Monte Carlo Computation and Bayesian Model Determination." *Biometrika* 82, no. 4 (1995): 711–732.
- Hoeting, Jennifer A., David Madigan, Adrian E. Raftery, and Chris T. Volinsky. "Bayesian Model Averaging: A Tutorial." *Statistical Science* 14, no. 4 (1999): 382–417.
- Kass, Robert E., and Adrian E. Raftery. "Bayes Factors." *Journal of the American Statistical Association* 90, no. 430 (1995): 773–795.
- Kuang, D., B. Nielsen, and J. P. Nielsen. "Identification of the Age–Period–Cohort Model and the Extended Chain-Ladder Model." *Biometrika* 95, no. 4 (2008): 979–986.
- Lange, Kenneth L., Roderick J. A. Little, and Jeremy M. G. Taylor. "Robust Statistical Modeling Using the t Distribution." *Journal of the American Statistical Association* 84, no. 408 (1989): 881–896.
- Lipton, Barry, John Robertson, and Katy Porter. "Workers Compensation Temporary Total Disability Indemnity Benefit Duration—2013 Update." Boca Raton, FL: National Council on Compensation Insurance, 2013.
- Mack, Thomas. "Distribution-Free Calculation of the Standard Error of Chain Ladder Reserve Estimates." *ASTIN Bulletin* 23, no. 2 (1993): 213–225.

Meyers, Glenn G. *Stochastic Loss Reserving Using Bayesian MCMC Models*. CAS Monograph Series 1. Arlington, VA: Casualty Actuarial Society, 2015.

Meyers, Glenn G., and Peng Shi. "The Retrospective Testing of Stochastic Loss Reserve Models." *Casualty Actuarial Society E-Forum* (Summer 2011): 1–37.

Schmid, Frank A. "The Workers Compensation Tails." *Variance* 6, no. 1 (2012): 48–77.

Schmid, Frank A., and Nathan Lord. "The Impact of Physician Fee Schedule Changes in Workers Compensation: Evidence From 31 States." Boca Raton, FL: National Council on Compensation Insurance, 2013.

Scott, James G., and James O. Berger. "Bayes and Empirical-Bayes Multiplicity Adjustment in the Variable-Selection Problem." *Annals of Statistics* 38, no. 5 (2010): 2587–2619.

Verrall, Richard J., and Mario V. Wüthrich. "Reversible Jump Markov Chain Monte Carlo Method for Parameter Reduction in Claims Reserving." *North American Actuarial Journal* 16, no. 2 (2012): 240–259.

Zanella, Giacomo. "Informed Proposals for Local MCMC in Discrete Spaces." *Journal of the American Statistical Association* 115, no. 530 (2020): 852–865.

Zellner, Arnold. "Bayesian and Non-Bayesian Analysis of the Log-Normal Distribution and Log-Normal Regression." *Journal of the American Statistical Association* 66, no. 334 (1971): 327–330.