Optimal Relativities in a Driving-Behavior Bonus–Malus System under Cluster-Aware Transition Rules

Mohamed Hanafy Kotb Ibrahim

Joint work with:

Eric C.K. Cheung, Andrés Villegas Ramirez, and Jae-Kyung Woo

School of Risk and Actuarial Studies

University of New South Wales

IDSC 2026 · Hannover · 9–10 June 2026

1 Background

Motivation: pricing on what we can actually *see*

- Traditional motor pricing uses **demographic proxies**: age, gender, location.
- Telematics records **driving behavior directly**: a far richer signal than any proxy.
- The Bonus–Malus System (BMS) adjusts premiums **ex post** from claim history.

Standard BMS limitations

- Mostly **claim-count based**; some extensions add severity (Oh et al. 2022) or near-miss events (Simon et al. 2025; Yanez et al. 2025); **behavioral clusters remain unused**.
- BMS fails as a risk classifier; bonus hunger; extreme malus tail.

Can we redesign the BMS so that behavior reshapes the entire stationary distribution?

What's known and what's missing

Telematics + insurance literature.

- Pay-As-You-Drive / Pay-How-You-Drive improve loss prediction
  (Ayuso et al. 2019; Chan et al. 2025a; Chan et al. 2025b; Duval et al. 2024; Gao et al. 2022)
- Frequency–severity dependence matters
  (Ahn et al. 2022; Cheung et al. 2021; Oh et al. 2022)
- Behavioral clustering captures stable risk types
  (Deng et al. 2022; Mantouka et al. 2019; Shirmohammadi et al. 2019; Yao et al. 2021)

The gap.

- Existing work uses telematics to improve **loss prediction** (the base premium).
- BMS *transitions* still depend on **claim counts**, **severity** (Oh et al. 2022), or **near-miss events** (Simon et al. 2025).
- **Behavioral clusters as a BMS transition channel remain unexplored.**

This paper: *the Driving-Behavior Bonus–Malus System (DBBMS).*

Data: a public telematics dataset

Source. Synthetic telematics dataset of So et al. (2021): a public benchmark in telematics-pricing research.

> $P(Y_i = 0) = \mathbf{96.3}\%$: *extreme zero-inflation motivates the Hurdle–Tweedie specification.*

Variable group	Count
*Traditional rating variables (**x**)*	
Demographics & vehicle (numeric)	7
Sex, marital, region, car use	4
Subtotal	**11**
*Telematics signals (**z**) — K-means clustering input*	
Exposure (% driven, miles, days/wk)	3
Day-of-week, period, rush-hour share	11
Trip-duration buckets	3
Acceleration / braking events	12
Turn-signal intensity, left / right	10
Subtotal	**39**
Claim outcomes	
NB_Claim (count), AMT_Claim (loss)	2

The DBBMS framework

Framework overview

The DBBMS: a two-stage plug-in framework.

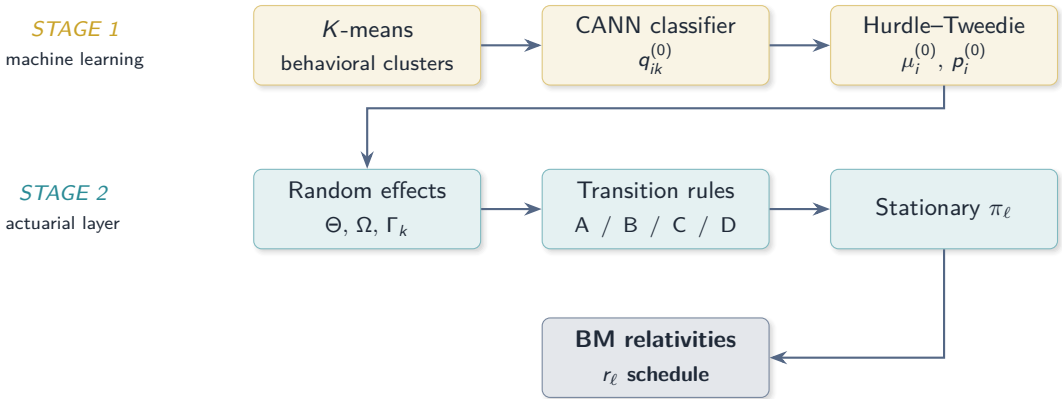

Stage 1 supplies fitted baselines; Stage 2 absorbs residual heterogeneity *without* re-estimating Stage 1.

Stage 1a: Behavioral clustering ($K = 3$)

K-means on telematics features;
$K = 3$ chosen by Elbow + Silhouette.

$G_i \in \{k^S, k^M, k^A\}$ denotes the behavioral cluster of policyholder i.

- k^S **Safe** (76.9%)
- k^M **Moderate** (10.3%)
- k^A **Aggressive** (12.8%)

Two minority clusters capture qualitatively different dimensions of aggressive driving –

- k^M (Moderate): hard braking, sharp acceleration.
- k^A (Aggressive): aggressive turns (left + right).

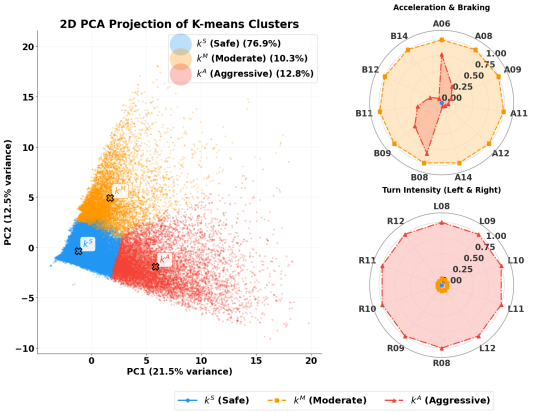

Principal component analysis (PCA) projection + radar profiles

Gated feature selection (GFS)

Architecture of one Combined Actuarial Neural Network (CANN) head

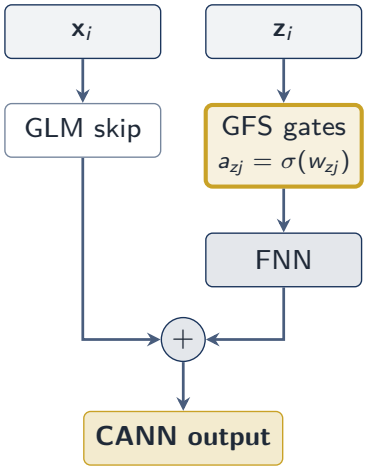

Telematics passes through GFS gates; covariates skip via Generalized Linear Model (GLM); outputs

Stage 1b: CANN classifier & Hurdle–Tweedie loss model

CANN cluster classifier: Multinomial Logistic Regression (MLR) baseline + FNN residual branch with GFS.

- GLM alone: 76.8% accuracy but **collapses** to k^S: zero recall on minorities.
- CANN: **99.8% accuracy** with only 8 GFS-selected features.

Loss model. Hurdle–Tweedie specification: Bernoulli zero-loss probability $p_i = \mathbb{P}(Y_i = 0)$ plus a zero-truncated Tweedie conditional on $Y_i > 0$.

Model	RMSE	MAE	Pred / Act
TwCANN (Tweedie CANN)	1,438	209	**0.517**
Hurdle–Tweedie CANN	**1,293** (↓ 10%)	**197** (↓ 6%)	**0.939** (↑ 82%)

Stage 2: Random effects and baseline preservation

Idea. Stage 1 gives the fitted baseline. Stage 2 adds residual heterogeneity, but uses centering terms so the baseline is preserved componentwise.

$$\tilde{\mu}_i^{\mathrm{Tw}} = \underbrace{\mu_i^{(0)}}_{\text{Stage-1 baseline}} \cdot \exp\Big(\underbrace{\Theta_i}_{\text{random effect}} - \underbrace{\tfrac{1}{2}\Sigma_{11}}_{\text{centering}} \Big)$$

Same idea on the other scales:

$$\tilde{p}_i = \mathrm{logit}^{-1}\big[\,\mathrm{logit}(p_i^{(0)}) + \Omega_i - c_i^{(0)}\,\big]$$

$$P(G_i = k \mid \cdot) = \frac{\exp\big(s_{ik} + \Gamma_{ik} - c_{ik}^{(g)}\big)}{\sum_{m=1}^{K} \exp\big(s_{im} + \Gamma_{im} - c_{im}^{(g)}\big)}$$

Random effects capture residual heterogeneity, while centering prevents Stage 2 from shifting the Stage 1 baselines.

Stage 2: The relativity refinement chain

Progressive refinement of the relativity.

$$r_\ell^{\text{cred}} \xrightarrow{+\Lambda^2\text{-weighting}} \tilde{r}_\ell \xrightarrow{+\,\mathfrak{m}\text{ equilibrium}} r_\ell^{*} \xrightarrow{+\text{linearity, monotonicity}} r_\ell^{\ddagger}$$

- **Credibility** $r_\ell^{\text{cred}} = \mathbb{E}[\Psi \mid L = \ell]$: the benchmark.
- **Weighted** $\tilde{r}_\ell = A_\ell / B_\ell$: weighted by Λ^2.
- **Optimal** $r_\ell^{*} = \tilde{r}_\ell - \mathfrak{m}/B_\ell$: enforces $\mathbb{E}[r(L)] = 1$.
- **Commercial** $r_\ell^{\ddagger} = a^{\ddagger} + b^{\ddagger}\ell$: linear, monotone, positive.

Notation: L random BM level, ℓ its realization; Λ baseline expected loss; Ψ latent-risk ratio;
$A_\ell = \mathbb{E}[\Lambda^2\Psi \mid L{=}\ell]$, $B_\ell = \mathbb{E}[\Lambda^2 \mid L{=}\ell]$; $\mathfrak{m}$ equilibrium multiplier; $a^{\ddagger}, b^{\ddagger}$ schedule intercept and slope.

Each step adds one constraint; the final $r_\ell^{\ddagger}$ is the commercially admissible schedule.

Four transition rules

Rule A	Classic claim-frequency BMS
	$\ell_{i,t+1} = \ell_{it} + \mathbb{I}(D_{it}{=}1) - \mathbb{I}(D_{it}{=}0)$
Rule B	Behavior-augmented (cluster-specific malus / bonus)
	$\ell_{i,t+1} = \ell_{it} + \mathbb{I}(D_{it}{=}1)\, m_{G_i} - \mathbb{I}(D_{it}{=}0)\, b_{G_i}$
Rule C	Severity-tiered (classes $\mathcal{S}^1, \mathcal{S}^2, \mathcal{S}^3$)
	$\ell_{i,t+1} = \ell_{it} - \mathbb{I}(D_{it}{=}0)\, b + \sum_{j=1}^{3} \mathbb{I}(Y_{it} \in \mathcal{S}^j)\, m_j$
Rule D	**Cluster $\times$ severity (full integration)**
	$\ell_{i,t+1} = \ell_{it} - \mathbb{I}(D_{it}{=}0)\, b_{G_i} + \sum_{j=1}^{3} \mathbb{I}(Y_{it} \in \mathcal{S}^j)\, m_{j,G_i}$

After each update, levels are truncated to $\{0, \ldots, \nu - 1\}$ (ν BM levels): $\ell_{i,t+1} \leftarrow \min\{\nu - 1, \max\{0, \ell_{i,t+1}\}\}$.

3 Empirical results

Main result: the rule reshapes the pooled stationary distribution

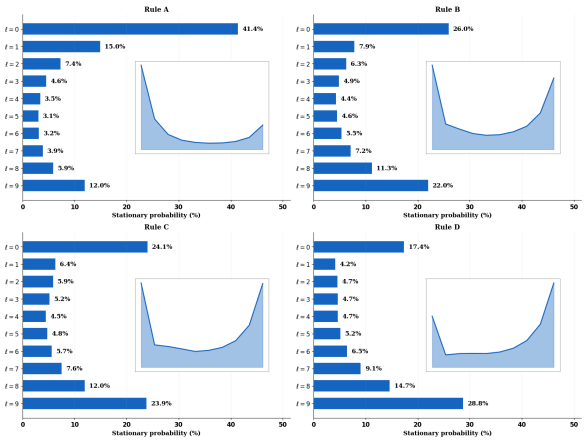

Rule	π_0	π_9	$\mathcal{S}_\pi^{\text{agg}}$
A (claims)	41.4%	12.0%	0.392
C (severity)	24.1%	23.9%	1.169
B (cluster)	26.0%	22.1%	1.020
D (full)	**17.4%**	**28.8%**	**1.801**

π_ℓ = long-run share of policyholders at BM level ℓ;

$\mathcal{S}_\pi^{\text{agg}} := \dfrac{\sum_{\ell \geq \lceil \nu/2 \rceil} \pi_\ell}{\sum_{\ell < \lceil \nu/2 \rceil} \pi_\ell}$ (malus-half / bonus-half mass).

$\mathcal{S}_\pi^{\text{agg}} > 1$ = malus-loaded regime.

> *Rule D drives a malus-loaded regime at **typical motor frequency**.*

Cluster Visibility

Rule	$\mathcal{S}_\pi^{\text{agg},k^S}$	$\mathcal{S}_\pi^{\text{agg},k^M}$	$\mathcal{S}_\pi^{\text{agg},k^A}$	Spread
Cluster-blind rules				
A (claims)	0.359	0.408	0.410	0.051
C (severity)	1.013	1.184	1.334	0.321
Cluster-aware rules				
B (cluster)	0.359	1.195	**2.402**	**2.043**
D (full)	**1.022**	**1.958**	**3.193**	**2.171**

Spread $= \mathcal{S}_\pi^{\text{agg},k^A} - \mathcal{S}_\pi^{\text{agg},k^S}$. All rules at typical motor frequency $\bar{\lambda} \approx 0.04$.

Per-cluster relativities under Rule D

Cluster	$r^{\ddagger}_{\ell=0}$	$r^{\ddagger}_{\ell=4}$	$r^{\ddagger}_{\ell=9}$	$\sum_\ell \pi^k_\ell \, r^{\ddagger,k}_\ell$
k^S (Safe)	**0.653**	0.964	1.352	1.000
k^M (Moderate)	0.438	0.833	1.328	1.000
k^A (Aggressive)	**0.364**	0.761	1.258	1.000
Pooled	0.483	0.858	1.326	1.000

Constrained linear schedule $r^{\ddagger,k}_\ell = a^{\ddagger,k} + b^{\ddagger,k}\,\ell$, equilibrium imposed within each cluster.

*At the best level $\ell = 0$, the Safe cluster pays **0.653** while the Aggressive cluster pays only **0.364**. Reading per-level relativities across clusters suggests Safe drivers are penalized more than Aggressive drivers but this misreads the construction.*

Why this matters: BMS as a true risk classifier

Right-skewed, no cluster separation

Large mass at the best bonus level.

- Weak classifier: a large share sit at the best level on the same discount.
- **Bonus hunger**: suppress small claims.
- Equilibrium forces extreme malus on a tiny tail.

Sharp cluster separation (Rule D, $\mathcal{S}_{\pi}^{\text{agg}} = 1.80$)

Three cluster tiers with distinct long-run positions.

- **Clear separation** of Safe, Moderate, Aggressive drivers.
- Bonus-hunger pathology mitigated.
- Moderate maluses spread across the portfolio.

The transition-rule channel alone controls the stationary shape, without changing the loss model or the relativity formula.

4 Conclusion

Conclusion

- **Telematics should enter the BM transition rule, not only the base premium.**
- **Cluster-aware rules make behavioral risk visible in the long run.**
- **Relativity comparisons must use expected premium, not one BM level.**

References i

Ahn, J. Y., E. C. K. Cheung, R. Oh, and J.-K. Woo (2022). **Optimal relativities in a modified Bonus–Malus system with long memory transition rules and frequency–severity dependence.** *Variance* 15(2). doi:10.66573/001c.46665.

Ayuso, M., M. Guillen, and J. P. Nielsen (2019). **Improving automobile insurance ratemaking using telematics: incorporating mileage and driver behaviour data.** *Transportation* 46(3), 735–752. doi:10.1007/s11116-018-9890-7.

Chan, I. W., A. L. Badescu, and X. S. Lin (2025a). **Assessing driving risk through unsupervised detection of anomalies in telematics time series data.** *ASTIN Bulletin* 55(2), 205–241. doi:10.1017/asb.2025.14.

Chan, I. W., S. C. Tseung, A. L. Badescu, and X. S. Lin (2025b). **Data mining of telematics data: Unveiling the hidden patterns in driving behavior.** *North American Actuarial Journal* 29(2), 275–309. doi:10.1080/10920277.2024.2376816.

Cheung, E. C., W. Ni, R. Oh, and J.-K. Woo (2021). **Bayesian credibility under a bivariate prior on the frequency and the severity of claims.** *Insurance: Mathematics and Economics* 100, 274–295. doi:10.1016/j.insmatheco.2021.06.003.

References ii

- Deng, Z., D. Chu, C. Wu, S. Liu, C. Sun, T. Liu, and D. Cao (2022). **A probabilistic model for driving-style-recognition-enabled driver steering behaviors.** *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 52(3), 1838–1851. doi:10.1109/TSMC.2020.3037229.
- Duval, F., J.-P. Boucher, and M. Pigeon (2024). **Telematics combined actuarial neural networks for cross-sectional and longitudinal claim count data.** *ASTIN Bulletin* 54(2), 239–262. doi:10.1017/asb.2024.4.
- Gao, G., S. Meng, and M. V. Wüthrich (2022). **What can we learn from telematics car driving data: A survey.** *Insurance: Mathematics and Economics* 104, 185–199. doi:10.1016/j.insmatheco.2022.02.004.
- Mantouka, E. G., E. N. Barmpounakis, and E. I. Vlahogianni (2019). **Identifying driving safety profiles from smartphone data using unsupervised learning.** *Safety Science* 119, 84–90. doi:10.1016/j.ssci.2019.01.025.
- Oh, R., J. H. Kim, and J. Y. Ahn (2022). **Designing a bonus–malus system reflecting the claim size under the dependent frequency–severity model.** *Probability in the Engineering and Informational Sciences* 36(4), 963–987. doi:10.1017/S0269964821000188.
- Shirmohammadi, H., F. Hadadi, and M. Saeedian (2019). **Clustering analysis of drivers based on behavioral characteristics regarding road safety.** *International Journal of Civil Engineering* 17(8), 1327–1340. doi:10.1007/s40999-018-00390-2.

References iii

- Simon, P.-A., J. Trufin, and M. Denuit (2025). **Bivariate Poisson credibility model and bonus–malus scale for claim and near-claim events.** *North American Actuarial Journal* 29(1), 74–93. doi:10.1080/10920277.2023.2293210.
- So, B., J.-P. Boucher, and E. A. Valdez (2021). **Synthetic dataset generation of driver telematics.** *Risks* 9(4), 58. doi:10.3390/risks9040058.
- Yanez, J. S., M. Guillén, and J. P. Nielsen (2025). **Weekly dynamic motor insurance ratemaking with a telematics signals bonus–malus score.** *ASTIN Bulletin* 55(1), 1–28. doi:10.1017/asb.2024.30.
- Yao, Y., X. Zhao, Y. Wu, Y. Zhang, and J. Rong (2021). **Clustering driver behavior using dynamic time warping and hidden Markov model.** *Journal of Intelligent Transportation Systems* 25(3), 249–262. doi:10.1080/15472450.2019.1646132.

Thank you.

Questions & discussion

mohamed_hanafy.ibrahim@student.unsw.edu.au

Data source. So et al. (2021): synthetic dataset calibrated against a real Canadian motor portfolio.

Preprocessing.

- Drop policies with any telematics variable above the 98th percentile.
- Drop positive-count / zero-loss inconsistencies.
- Standardize numerics on training-set statistics only (no leakage).

BMS configuration.

Quantity	Value
Mean claim frequency $\bar{\lambda}$	≈ 0.04
Severity thresholds (ξ_1, ξ_2)	(1,037, 3,097)
ν (BM levels) / ε (tol.)	10 / 0.01

ξ_1, ξ_2 are the 33rd and 67th percentiles of the positive-loss distribution.

Appendix • Decile calibration: Hurdle–Tweedie vs. pure Tweedie

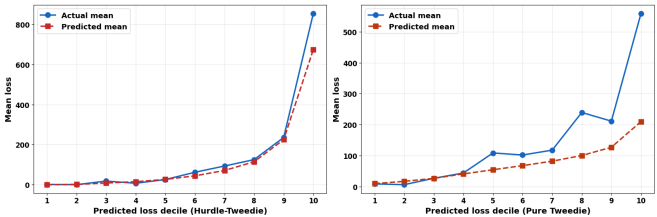

Mean actual vs. mean predicted aggregate loss within each decile of predicted loss, on the test set. Left: Hurdle–Tweedie CANN. Right: pure TwCANN.

- Hurdle–Tweedie tracks actual losses closely across **deciles 1–9**.
- Residual top-decile under-prediction $\approx$ **20%**.
- Pure Tweedie under-predicts from the **fifth decile** onward, reaching $\approx$ **60%** in the top decile.

The atom-at-zero limitation of the pure Tweedie shows up precisely where losses are largest.

Latent-risk adjustment. For each policy i,

$$\Lambda_i = \frac{(1 - p_i^{(0)})\,\mu_i^{(0)}}{1 - \exp(-w_i^{(0)})}, \qquad \Psi_i(\Theta_i, \Omega_i) = \frac{\tilde{\Lambda}_i(\Theta_i, \Omega_i)}{\Lambda_i}.$$

$w_i^{(0)}$ is the Tweedie zero-mass exponent; $\tilde{\Lambda}_i$ is the random-effect-adjusted target.

Stationary moments at level ℓ: $A_\ell := \mathbb{E}[\Lambda^2\Psi \mid L = \ell]$, $B_\ell := \mathbb{E}[\Lambda^2 \mid L = \ell]$.

Four relativities, in closed form:

$$r_\ell^{\text{cred}} = \mathbb{E}[\Psi \mid L = \ell] \qquad \text{(Bühlmann benchmark)}$$

$$\tilde{r}_\ell = A_\ell / B_\ell \qquad (\Lambda^2\text{-weighted})$$

$$r_\ell^{*} = \tilde{r}_\ell \;-\; \mathfrak{m}/B_\ell, \quad \mathfrak{m} = \frac{\sum_\ell \pi_\ell\, \tilde{r}_\ell - 1}{\sum_\ell \pi_\ell / B_\ell} \qquad \text{(optimal, equilibrium)}$$

$$r_\ell^{\ddagger} = a^{\ddagger} + b^{\ddagger}\ell \qquad \text{(commercial: monotone, positive)}$$

Lagrangian derivation in Appendix B.1 of the paper; constrained linear projection in Appendix B.3.

Weighted moments. $M_k := \mathbb{E}[\Lambda^2 L^k]$ for $k = 0, 1, 2$; $W_0 := \mathbb{E}[\Lambda^2 \Psi]$, $W_1 := \mathbb{E}[\Lambda^2 \Psi L]$; $\bar{L} := \mathbb{E}[L]$.

Unconstrained Λ^2-weighted least squares: minimises $\mathbb{E}[\Lambda^2(\Psi - (a + bL))^2]$:

$$b_{\mathrm{lin}} = \frac{W_1 - M_1 W_0 / M_0}{M_2 - M_1^2 / M_0}, \qquad a_{\mathrm{lin}} = \frac{W_0 - b_{\mathrm{lin}} M_1}{M_0}.$$

Equilibrium-constrained projection ($\mathbb{E}[r(L)] = 1$, $b \geq 0$, $a \geq \varepsilon$):

$$b^{\ddagger} = \min\left\{ \max\{0,\, b_{\mathrm{bal}}\},\ \tfrac{1-\varepsilon}{\bar{L}} \right\}, \qquad a^{\ddagger} = 1 - b^{\ddagger}\, \bar{L},$$

where $b_{\mathrm{bal}} = \dfrac{\mathbb{E}[\Lambda^2(\Psi - 1)(L - \bar{L})]}{\mathbb{E}[\Lambda^2(L - \bar{L})^2]}$.

> $r_\ell^{\ddagger}$ *projects* b_{bal} *onto* $[0, (1-\varepsilon)/\bar{L}]$*: monotonicity, positivity, and equilibrium all at once.*

Two complementary diagnostics agree on $K = 3$.

- **Elbow method:** within-cluster sum of squares flattens at $K = 3$; marginal reduction from $K = 3$ to $K = 4$ is small.
- **Average Silhouette Width:** $\mathrm{ASW}(K) = \frac{1}{n}\sum_{i=1}^{n} s_i$, $s_i \in [-1, 1]$ peaks at $K = 3$ and decays for $K \geq 4$.

Why this matters for the BMS. A coarser partition (smaller K) preserves stable, well-separated risk types; a finer partition risks low-occupancy clusters with less stable membership, which would propagate to noisy cluster-conditional transition probabilities.

Estimation. Conditional plug-in observed-data criterion given fitted Stage 1 baselines:

$$\mathcal{L}_{\mathrm{obs}}(\boldsymbol{\Sigma}) \;=\; \sum_{i=1}^{n} \log \int_{\mathbb{R}^{K+1}} \mathcal{I}_i\big(\Theta_i, \Omega_i, \Gamma_{i,2:K}\big)\, f_{\boldsymbol{\Sigma}}\big(\Theta_i, \Omega_i, \Gamma_{i,2:K}\big)\, \mathrm{d}\Theta_i\, \mathrm{d}\Omega_i\, \mathrm{d}\Gamma_{i,2:K},$$

where $\mathcal{I}_i$ is the integrand assembled from the Hurdle–Tweedie and cluster components, and $f_{\boldsymbol{\Sigma}}$ is the joint normal density on $(\Theta_i, \Omega_i, \Gamma_{i,2:K})$.
Integrals over the random effects evaluated by the Laplace approximation at the modal point (Appendix A.3).

Estimated random-effect parameters.

Component	$\hat{\sigma}^2$	$\hat{\sigma}$
σ_Θ (Tweedie mean)	1.095	1.046
σ_Ω (hurdle binary)	0.216	0.465
σ_{Γ_2} (cluster k^M)	< 0.001	0.013
σ_{Γ_3} (cluster k^A)	< 0.001	0.018

Mixture-integral form of $A_\ell = \mathbb{E}[\Lambda^2 \Psi \mid L = \ell]$ (and analogously B_ℓ):

$$A_\ell \;=\; \frac{1}{\pi_\ell} \sum_{\mathfrak{g}=1}^{\mathcal{G}} \varpi_{\mathfrak{g}} \int \sum_{k=1}^{K} \Lambda_{\mathfrak{g}}^2 \, \Psi_{\mathfrak{g}}(\theta, \omega) \, \pi_\ell^{(\mathfrak{g})}\big(\Xi_{\mathfrak{g}}(\theta, \omega), k\big) \, p_k^{(\mathfrak{g})}(\boldsymbol{\gamma}_{2:K}) \, f_{\boldsymbol{\Sigma}} \, \mathrm{d}\theta \, \mathrm{d}\omega \, \mathrm{d}\boldsymbol{\gamma}_{2:K}.$$

Two nested loops:

- **Outer:** Laplace approximation at the modal point of the random-effect distribution over $(\theta, \omega, \boldsymbol{\gamma}_{2:K})$.
- **Inner:** stationary distribution $\pi^{(\mathfrak{g})}$ of a $\nu \times \nu$ Markov chain: closed-form via the dominant left eigenvector.

> The plug-in structure means Stage 1 (the expensive neural network) is fit once and reused for every Stage 2 transition-rule evaluation.