https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

A traceable grid-based framework for earthquake parametric insurance pricing: methodology and case studies in Indonesia

Hartati [1,3], Adhitya Ronnie Effendie[1], Nanang Susyanto[1], and Wiwit Suryanto[2]

[1]Department of Mathematics, Faculty of Mathematics and Natural Sciences, Universitas Gadjah Mada, Yogyakarta 55281, Indonesia

[2]Department of Physics, Faculty of Mathematics and Natural Sciences, Universitas Gadjah Mada, Yogyakarta 55281, Indonesia

[3]Mathematics Study Program, Faculty of Sciences and Technology, Open University, Tangerang Selatan 15437, Indonesia

Correspondence: Adhitya Ronnie Effendie (adhityaronnie@ugm.ac.id)

Abstract. We propose $\mathcal{F}_g$, a portable and end-to-end traceable framework for pricing earthquake parametric insurance contracts on a fixed grid, and apply it to Indonesia as a case study. Three persistent gaps motivate the framework. National PSHA delivers a single value per return period, not the full annual distribution that contract calibration requires; existing risk metrics are reported at portfolio level and lose grid-level traceability; and basis risk is evaluated post hoc rather than as a design feedback. $\mathcal{F}_g$ closes these gaps by combining four established components on a single $1^\circ \times 1^\circ$ spatial unit: (i) annual hazard via Gutenberg–Richter recurrence with a Boore–Atkinson 2008 baseline GMPE, (ii) a Grid Loss Table from Monte Carlo simulation ($N = 20\,000$), (iii) parametric contract calibration via local AEP inversion, and (iv) a basis-risk vector fed back into the contract before the premium is set. The resulting premium Π_g is a deterministic function of just four catalogue inputs $\Theta_g = (n_b, \hat{\lambda}_b, \hat{b}_b, E_g)$. Applying $\mathcal{F}_g$ to the USGS catalogue for Indonesia (1925–2025; 84 774 events; 623 active grids; the Gardner–Knopoff dispersion index drops from 111.4 to 4.2 after declustering) produces national maps of the seismicity parameters and average annual loss, plus end-to-end pricing on five representative grids spanning the national variation. The actuarial loading ratio $\Pi_g/\mathbb{E}[Y_g]$ falls in a 5.87–6.13 band across the four high- to medium-hazard grids—a property that, as we show analytically, follows directly from the AEP-inversion calibration combined with the fixed loading scheme. Two independent sensitivity experiments on grid G1 confirm this structural prediction numerically. Substituting the BC Hydro subduction GMPE for the BA08 baseline shifts the loading ratio by only 1.7 %, and substituting either the Reasenberg or the Zaliapin–Ben-Zion declustering algorithm for the Gardner–Knopoff baseline shifts it by less than 1 %—three substitutions in total across the two experiments, all of them moving the absolute premium by 22–94 % but leaving the loading ratio essentially unchanged. The very-low-hazard reference grid never reaches its parametric trigger, yielding an explicit segmentation criterion that separates regions suitable for parametric protection from those better served by indemnity coverage. Every premium is auditable end-to-end from catalogue to tariff, and the four-module structure transfers to any seismic region by substituting the catalogue, GMPE, vulnerability curve, and exposure database.

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

1 Introduction

Indonesia sits at the convergence of four tectonic plates and accumulates large-magnitude earthquakes at a rate few other countries match (Fig. 1). When such an event strikes, the speed of post-disaster liquidity is decisive for households and small enterprises, yet conventional indemnity insurance—which requires post-event damage assessment—typically takes months to settle (World Bank, 2017). Parametric insurance, in which payouts are triggered by a measurable index, removes this delay (Swiss Re Institute, 2016; Franco, 2010) at the cost of basis risk: the residual gap between the index and the actual loss (Cummins, 2012; Lin and Shaw, 2008).

The seismic-hazard tools needed to design such contracts already exist. Indonesia has a national PSHA (Irsyam et al., 2020), the Global Earthquake Model provides a global comparator (Pagani et al., 2023), and OpenQuake supports event-based risk pipelines at scale (Silva et al., 2020). Yet three operational limitations have prevented these tools from being directly used for parametric pricing. National PSHA returns a single PGA per return period, not the full annual distribution $F_{I_g}(x)$ that contract calibration requires. Recent national-scale risk studies—Salgado-Gálvez et al. (2024) for Central Asia, Ghafory-Ashtiany and Motamed (2024) for Iran—produce aggregate AAL at portfolio or oblast level, losing per-grid traceability to local seismicity. And the parametric-insurance literature treats basis risk as a post-design diagnostic (Mousavi et al., 2023; Pai et al., 2022), not as a feedback signal that drives recalibration. Each gap is independently bridged in some part of the literature; none has been bridged simultaneously.

This paper addresses these gaps through $\mathcal{F}_g$, an integrated and traceable framework that operates on a fixed $1^\circ \times 1^\circ$ grid as a four-module pipeline: (i) the per-grid annual maximum index is stored as a full distribution; (ii) annual losses are tabulated per grid in a Grid Loss Table that replaces the conventional Event Loss Table; (iii) the contract attachment A_g^* and exhaustion X_g^* are calibrated by local AEP inversion to a probability target common across all grids; and (iv) a basis-risk vector is computed from paired (loss, payout) draws and fed back into the contract before the premium is set.

The contribution lies in the synthesis, not the components. The Poisson model, the Gutenberg–Richter law, the BA08 GMPE, and the actuarial premium formula are each well established; what we propose is their integration on a single spatial unit with end-to-end traceability. Every premium Π_g that $\mathcal{F}_g$ produces is a deterministic function of just four catalogue inputs, $\Theta_g = (n_b, \hat{\lambda}_b, \hat{b}_b, E_g)$—auditability that is essential for insurance regulators and disaster-risk-financing policy makers requiring a transparent basis for spatial tariff differentiation (Earthquake Index Insurance Initiative, 2016; OECD, 2021).

We validate $\mathcal{F}_g$ at two levels. At the national level, Modules 1–2 are run over all 623 active Indonesian grids to map the spatial structure of seismicity, exposure, and risk. At the case-study level, the full four-module chain is executed end-to-end on five grids selected to span the national variation in $(\hat{\lambda}, \hat{b}, E_g)$: two Sunda subduction grids, two collision/strike-slip grids, and one low-hazard reference. Two findings emerge. First, the actuarial loading ratio $\Pi_g/\mathbb{E}[Y_g]$ falls in a narrow 5.87–6.13 band across the four high- to medium-hazard grids. We show in Sect. 5.2 that this is a mathematical consequence of the AEP-inversion calibration combined with the fixed loading scheme rather than an emergent empirical finding; we further confirm this structural prediction with two independent quantitative sensitivity tests on grid G1 (Sect. 5.5). Replacing BA08 with the BC Hydro subduction GMPE shifts the loading ratio by only 1.7 % even though the absolute premium falls by 94 %;

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

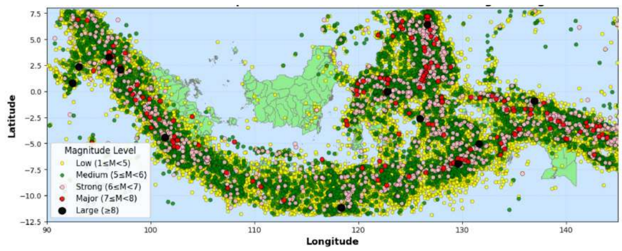

Figure 1. Epicentre map of earthquakes in Indonesia, 1925–2025, by magnitude class. Source: USGS catalogue, processed by the authors.

replacing Gardner–Knopoff declustering with either Reasenberg 1985 or Zaliapin–Ben-Zion 2013 shifts the loading ratio by less than 1 % even though the absolute premium changes by 22–26 %. Second, the low-hazard reference grid never reaches its parametric trigger, producing an explicit segmentation criterion: at this hazard level, parametric protection is not the appropriate instrument.

$\mathcal{F}_g$ is portable. Its four-module structure is not tied to any country-specific assumption; substituting the catalogue, regional GMPE, vulnerability curve, and exposure database adapts it to the Philippines, Chile, Japan, or Türkiye. Section 2 describes the data and pre-processing. Section 3 formalises the framework. Section 4 reports the national context. Section 5 presents the end-to-end results on five representative grids together with their interpretation. Section 6 concludes.

2 Data and pre-processing

$\mathcal{F}_g$ is built on the USGS earthquake catalogue for Indonesia, 1925–2025. After curation (removal of duplicates and entries with missing magnitude), 84 774 events remain in the region $90.03°$–$144.999°$ E and $-12.49°$–$8.00°$ N (Table 1). The catalogue mixes magnitude scales, so we homogenise to M_w using direct values where available (8 066 events; 10 %), the regional empirical relations of Diantari et al. (2018) for the dominant M, m_b, M_L, $M_\mathrm{w}(m_B)$ scales (76 349 events; 90 %), and the global Scordilis (2006) conversions as fallback for M_s and M_D (359 events; $< 1\,\%$). The Hanks–Kanamori relation is used where only seismic moment is reported (Hanks and Kanamori, 1979). Homogenisation to M_w is essential because mixing scales biases both b-value and rate estimates (Aki, 1965; Utsu, 1999).

Although the curated catalogue spans 1925–2025, we restrict the rate and b-value estimation in $\mathcal{F}_g$ to the sub-period 1980–2024 ($T_\mathrm{obs} = 45$ years). Pre-1980 instrumentation in the Indonesian region is dominated by the WWSSN network with limited spatial density and uneven detection thresholds across the archipelago, so the apparent rates and magnitude distributions before 1980 reflect monitoring capability rather than tectonics. From the late 1970s onward, the maturation of digital broadband networks and the rapid expansion of regional and global stations produced a far more spatially proportional record (see Storchak et al., 2013, on the homogeneity of the post-1976 ISC catalogue, of which the USGS PDE is the operational successor). The 1980–2024 window therefore offers the most defensible compromise between sample size and temporal homogeneity for parameter estimation. The pre-1980 events remain in the catalogue for completeness checks and for visualising the long-term spatial pattern (Figs. 1 and 2), but do not enter Modules 1–4.

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

Table 1. Curated USGS earthquake catalogue for Indonesia, 1925–2025.

Parameter	Value
Number of events	84 774
Period	1925–2025 (100 years)
Latitude / longitude range	$-12.49°$–$8.00°$ / $90.03°$–$144.999°$
Median depth	35.00 km
M_w range (mean $\pm$ SD)	3.13–9.10 (4.57 $\pm$ 0.46)

The completeness magnitude is estimated by the Maximum Curvature method (Wiemer and Wyss, 2000) with the conservative correction $\delta = 0.2$ of Woessner and Wiemer (2005). For the Sumatra subcatalogue we obtain $M_\mathrm{c} = 4.23$ (Fig. 3); we adopt $M_\mathrm{c} = 4.5$ as a conservative national working threshold. We caution that M_c varies across Indonesia: Java and Sumatra benefit from denser instrumentation and likely admit a lower local M_c, whereas Papua, Maluku, and the Banda Sea remain comparatively under-sampled and may have a locally higher M_c approaching 4.7–5.0. Adopting a single national threshold $M_\mathrm{c} = 4.5$ is conservative for the well-monitored regions (it discards some genuine events that would be detectable locally) and may be slightly optimistic for the under-monitored regions (where the true M_c might exceed 4.5 in some grids). We discuss the propagation of this assumption in Sect. 5.5; a region-by-region M_c map is the natural follow-up.

The study region is partitioned into $1° \times 1°$ cells. The choice of resolution is a trade-off: a finer grid (e.g. $0.5° \times 0.5°$) would resolve sub-provincial heterogeneity but would leave only $\sim 50\,\%$ of cells with $n_{b,g} \geq 25$ events over our 45-year window, dropping the valid grid count from 623 to roughly 250 and producing wider confidence intervals on $\hat{\lambda}_{b,g}$ and $\hat{b}_{b,g}$. A coarser grid ($2° \times 2°$) would oversmooth the Sunda Arc, where the hazard gradient across ~ 100 km is steep. The $1°$ resolution is a compromise that retains national coverage while keeping the per-grid sample size adequate for the Aki–Utsu MLE; it matches the resolution used in operational Indonesian seismicity studies and in the GEM exposure database. Sensitivity to grid size is a useful follow-up but does not affect the framework structure: $\mathcal{F}_g$ is independent of the chosen partition, and the same four-module pipeline can be re-run on a finer grid wherever the data permit. The event-density heatmap (Fig. 2) shows expected concentrations along the Sunda Arc, the Banda Arc, and the Sulawesi–Maluku belt, with sparse activity in Kalimantan and inner Papua.

The Poisson assumption underlying our hazard module requires temporal independence between events, which is violated by aftershock sequences. We apply Gardner–Knopoff (1974, hereafter GK74) declustering with the classical magnitude-dependent space–time window. To verify that the resulting background catalogue approximates a Poisson process, we compute the dispersion index $\mathrm{DI}_g = \mathrm{Var}(N_{y,g})/\mathbb{E}[N_{y,g}]$, which equals one for a Poisson process. The 1980–2024 catalogue gives $\mathrm{DI} = 111.4$ in the original series and $\mathrm{DI} = 4.2$ after declustering (Table 2)—a 96 % reduction. The aftershock sequence following the 2004 $M_\mathrm{w}\,9.1$ Sumatra–Andaman event dominates the original-catalogue overdispersion (Fig. 4); after declustering, annual counts stabilise in the range 208–352. The residual $\mathrm{DI} > 1$ is consistent with the literature for tectonically heterogeneous regional cat-

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

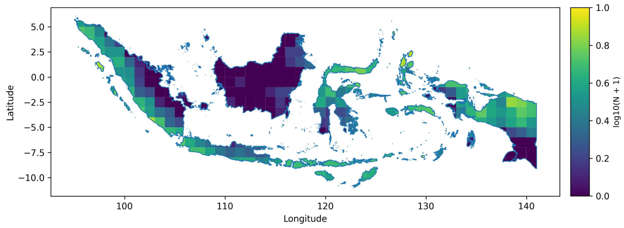

Figure 2. Event-density heatmap on a $1^\circ \times 1^\circ$ grid ($M_\mathrm{w} \geq 4.5$, 1925–2025). Colour shows $\log_{10}(N+1)$.

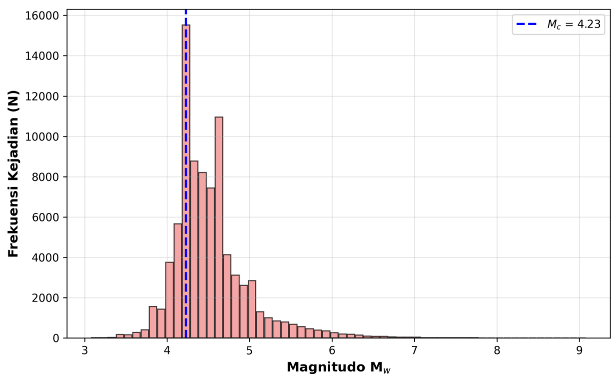

Figure 3. Frequency–magnitude distribution for the Sumatra subcatalogue with the completeness magnitude $M_\mathrm{c} = 4.23$ (vertical dashed line) from MAXC.

alogues at large spatial scale (Luen and Stark, 2012). We treat the Poisson model as a first-order approximation; sensitivity to alternative declustering algorithms is discussed in Sect. 5.5. All subsequent analyses use the declustered background catalogue, $M_\mathrm{w} \geq 4.5$, 1980–2024.

3 The framework $\mathcal{F}_g$

$\mathcal{F}_g$ is a four-module pipeline (Table 3). We describe each module in turn, then state the traceability proposition that links them.

3.1 Annual hazard

Background seismicity in grid g is modelled as a homogeneous Poisson process with annual rate $\lambda_{b,g}$, estimated by maximum likelihood as $\hat{\lambda}_{b,g} = n_{b,g}/T_\mathrm{obs}$ over $T_\mathrm{obs} = 45$ years. Magnitudes follow the Gutenberg–Richter law $\log_{10} N(M \geq m) = a - bm$, with the b-value estimated by the Aki–Utsu MLE,

$$\hat{b}_{b,g} = \frac{\log_{10} e}{\bar{M}_{b,g} - (M_\mathrm{c} - \Delta M/2)}, \qquad \Delta M = 0.1, \tag{1}$$

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

Table 2. Dispersion index before and after Gardner–Knopoff declustering ($M_\mathrm{w} \geq 4.5$, 1980–2024).

Catalogue	Mean / yr	Var / yr	DI
Original	763.07	84 991.29	111.4
Declustered (background)	287.24	1 196.42	4.2

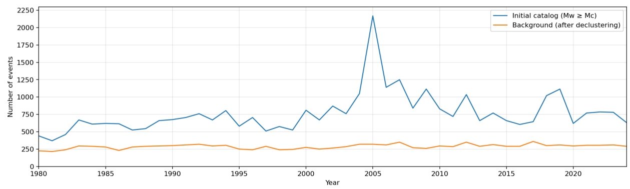

Figure 4. Annual event counts: original catalogue versus declustered background, 1980–2024. The 2005 aftershock peak following the 2004 Sumatra–Andaman event is largely removed by GK74.

where $\bar{M}_{b,g}$ is the sample mean of M_w in grid g over the declustered background catalogue, and $M_\mathrm{c} = 4.5$ is the national working threshold defined in Sect. 2. Grids with $n_{b,g} < 25$ do not satisfy the minimum sample size for a reliable Aki–Utsu MLE; for these grids we adopt the national-average $\hat{b} = 1.01$ obtained in our previous analysis of the same Indonesian USGS catalogue (Hartati et al., 2026), in which a Gutenberg–Richter regression on the country-wide frequency–magnitude distribution gave $a = 9.235$, $b = 1.01$ ($R^2 = 0.998$). This single-value imputation is conservative in the sense that it preserves the framework's national coverage at the cost of locally averaging out grid-specific deviations. The choice is also internally consistent with the present paper: the country-wide $b = 1.01$ falls inside the empirical range 0.57–1.53 (median 1.03) we observe directly on the 197 grids that satisfy $n_{b,g} \geq 25$ (Sect. 4). Sensitivity to the imputation rule is discussed in Sect. 5.5.

We adopt the Boore–Atkinson 2008 GMPE (Boore and Atkinson, 2008): $\ln \mathrm{PGA} = f(M_\mathrm{w}, R, V_{S30}; \theta) + \varepsilon$ with $\varepsilon \sim \mathcal{N}(0, \sigma^2)$. BA08 is calibrated primarily on shallow crustal earthquakes from active tectonic regions and is therefore not the optimal regression for the Sunda subduction interface and intraslab events that dominate Indonesian seismicity. Subduction-specific equations such as Zhao et al. (2006) for Japanese subduction sources, the BC Hydro model of Abrahamson et al. (2016) for global subduction, or the Indonesia-fitted regression of Megawati et al. (2003) for the Sumatra interface would in principle return median PGAs that differ from BA08 by 10–40 % at moderate magnitude and short distance. We adopt BA08 in this paper as a transparent, well-documented baseline rather than as the operationally optimal choice: it is the GMPE for which the Indonesian PSHA infrastructure is best documented and most easily reproduced (Irsyam et al., 2020; Pagani et al., 2023); the framework $\mathcal{F}_g$ is GMPE-agnostic in its structure (Eq. 2 accepts any $\Pr[\mathrm{PGA} > x \mid M, R]$); and a single-GMPE choice keeps the traceability chain (6) auditable in a way that a logic-tree weighted average does not. Substitution by a subduction-specific GMPE or by a logic tree of subduction equations is a one-line change in the OpenQuake configuration we use, leaving Modules 2–

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

Table 3. Summary of the four modules of $\mathcal{F}_g$. Each row lists the module's input, principal output, and the equation that defines the transformation. The output of each module is the input of the next.

Module	Input	Principal output	Defining Eq.
1. Annual hazard	$\Theta_g = (n_b, \hat{\lambda}_b, \hat{b}_b, E_g)$	$F_{I_g}(x) = \Pr(\mathrm{PGA}_{\max} \leq x)$	(2)
2. Grid Loss Table	F_{I_g}, E_g, vulnerability	GLT_g, AAL, VaR, TVaR, EP curve	—
3. Contract calibration	F_{I_g}, GLT_g, targets (p_A, p_X)	$(A_g^*, X_g^*, P_g^{\max})$, $Y_g(I)$	(3), (4)
4. Basis-risk feedback & premium	paired (L, Y) from GLT_g	BR_g vector, premium Π_g	(5)

4 unchanged. We quantify the propagation of the GMPE choice to the premium through a side-by-side BA08-vs-BC Hydro comparison on grid G1 in Sect. 5.5.4.

The site-condition parameter is held constant at $V_{S30} = 760\,\mathrm{m\,s^{-1}}$ across all grids in this paper, corresponding to a generic firm-rock reference. Indonesian site conditions in fact span from competent rock along Bukit Barisan and the inner Banda Arc to deep alluvium beneath Jakarta, Surabaya, and the northern Java coast; site amplification can shift PGA by a factor of 1.5–2 at low–moderate frequency. A grid-specific V_{S30} map (e.g., the USGS global V_{S30} slope-based proxy, or the SNI 1726:2019 site classification) is the natural next-generation input. As with the GMPE choice, V_{S30} enters $\mathcal{F}_g$ only through Module 1; Modules 2–4 are unaffected, and the overall traceability chain is preserved.

The annual hazard curve is the standard PSHA convolution (Cornell, 1968),

$$\lambda_{\mathrm{PGA},g}(x) = \sum_s \lambda_{b,s} \int_{M_\mathrm{c}}^{M_{\max}} \Pr[\mathrm{PGA} > x \mid M, R(s,g)]\, f_M(m)\,\mathrm{d}m, \tag{2}$$

where the sum runs over source cells s that contribute ground motion to the target grid g (in our default implementation, s enumerates all $1^\circ \times 1^\circ$ background cells within $R_\mathrm{hyp} \leq 300\,\mathrm{km}$ of the centroid of g); $\lambda_{b,s}$ is the annual rate of $M_\mathrm{w} \geq M_\mathrm{c}$ events in source s; $R(s,g)$ is the hypocentral distance from the centroid of s to the centroid of g; the integral runs over the truncated Gutenberg–Richter density $f_M(m)$ derived from $\hat{b}_{b,s}$ on $[M_\mathrm{c}, M_{\max}]$; and $M_{\max}$ is the source-zone upper bound (we use $M_{\max} = 9.0$ for the Sunda interface, consistent with the recurrence of the 2004 M_w 9.1 Sumatra–Andaman event). The annual maximum index has CDF $F_{I_g}(x) = \exp[-\lambda_{\mathrm{PGA},g}(x)]$. Storing F_{I_g} as a full distribution rather than a single point at one return period is what distinguishes $\mathcal{F}_g$ from conventional PSHA outputs and is the prerequisite for everything that follows.

3.2 Grid Loss Table and risk metrics

The annual loss at grid g in policy year y is $L_{g,y} = E_g \cdot \mathrm{LR}(I_{g,y})$, where $I_{g,y}$ is the annual maximum PGA at grid g in year y—drawn from F_{I_g} as defined in Sect. 3.1—and $\mathrm{LR}(\cdot)$ is a vulnerability function (we use the FEMA Hazus 5.1 baseline

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

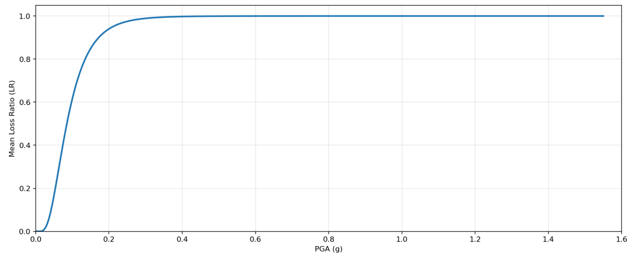

Figure 5. Baseline vulnerability function $\mathrm{LR}(\mathrm{PGA})$ adopted in Module 2 (FEMA Hazus 5.1, generic structural class).

shown in Fig. 5; Indonesia-specific calibration is discussed in Sect. 5.6). This formulation treats annual loss as a function of the dominant single event per grid-year and therefore omits cumulative damage from multiple sub-trigger events; the omission is conservative for the $\mathrm{TVaR}_{0.95}$-class years that drive the premium and is a known simplification of single-event annualisation in catastrophe modelling. Monte Carlo simulation with $N = 20\,000$ samples per grid produces the Grid Loss Table, $\mathrm{GLT}_g = \{(I_{g,y}, L_{g,y})\}_{y=1}^{N}$. Risk metrics—$\mathrm{AAL}$, VaR_α, TVaR_α, and the EP curve—are computed empirically. Unlike the conventional Event Loss Table, which aggregates losses across assets, the GLT preserves the per-grid–per-year structure that downstream modules require. This substitution is not cosmetic: the contract calibration in Sect. 3.3 requires the joint distribution of (I, L) at each grid, which the ELT structure cannot deliver.

3.3 Parametric contract calibration

The contract has three parameters: an attachment A_g (trigger), an exhaustion X_g (full payout), and a liability cap $P_g^{\max}$. We calibrate A_g and X_g by inverting the local hazard distribution to a common probability target across all grids,

$$A_g^* = F_{I_g}^{-1}(1 - p_A), \qquad X_g^* = F_{I_g}^{-1}(1 - p_X), \qquad p_X < p_A, \tag{3}$$

with $p_A = 0.10$ (10 %/yr trigger) and $p_X = 0.02$ (2 %/yr exhaustion). This delivers spatial protection equivalence: grids of different hazard receive numerically different thresholds but the same activation probability. Numerically, A_g^* and X_g^* coincide with the PGA at return periods $1/p_A$ and $1/p_X$, so internal consistency between the contract and the hazard summary is enforced by reading $\lambda_{\mathrm{PGA},g}$ once and reusing it. The liability cap is set from the loss distribution: $P_g^{\max} = \mathrm{VaR}_{0.95}(L_g)$. The piecewise-linear payout function is then

$$Y_g(I) = \begin{cases} 0, & I \le A_g^*, \\ P_g^{\max} \dfrac{I - A_g^*}{X_g^* - A_g^*}, & A_g^* < I < X_g^*, \\ P_g^{\max}, & I \ge X_g^*. \end{cases} \tag{4}$$

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

3.4 Basis risk and the risk-based premium

Define $D_{g,y} = L_{g,y} - Y_{g,y}$. The basis-risk vector is $\mathrm{BR}_g = (\mathrm{Bias}, \mathrm{MAE}, \rho(L,Y), \mathrm{EU}, p_{\mathrm{under}})$, where EU is the expected underpayment and $p_{\mathrm{under}} = \Pr(L > Y)$. Reporting a vector rather than a single scalar is essential: a single correlation, RMSE, or bias cannot distinguish three pathological cases—small bias with high p_{under}, large bias with low p_{under}, and small bias with low p_{under}—that the vector resolves. Critically, $L_{g,y}$ and $Y_{g,y}$ must be computed from the same draw $I_{g,y}$; computing them from independent draws produces spurious decoupling, inflates p_{over}, and obscures the true basis-risk profile. Our simulation enforces pairing by drawing $I_{g,y}$ once per grid-year and deriving both L and Y deterministically. If BR_g exceeds an acceptable threshold the contract is recalibrated before the premium is set. The risk-based premium is

$$\Pi_g = (1 + c_{\mathrm{exp}})\,\mathbb{E}[Y_g] + \gamma\big(\mathrm{TVaR}_\alpha(Y_g) - \mathbb{E}[Y_g]\big), \tag{5}$$

with default $c_{\mathrm{exp}} = 0.20$ (expense and margin), $\gamma = 0.30$ (tail loading), and $\alpha = 0.95$ (Klugman et al., 2012). The first term is the risk-loaded pure premium; the second penalises tail uncertainty.

3.5 Traceability

Equations (2)–(5) compose into a single chain,

$$\Theta_g = (n_b, \hat{\lambda}_b, \hat{b}_b, E_g) \longrightarrow F_{I_g} \longrightarrow L_g \longrightarrow Y_g \longrightarrow \Pi_g. \tag{6}$$

Two properties of (6) characterise the integration we propose. First, all transformations operate on the same spatial unit, so heterogeneity is preserved rather than aggregated away to portfolio level. Second, A_g^*, X_g^*, and $P_g^{\max}$ are calibrated from the same F_{I_g} that drives L_g, so the contract, the loss, and the premium are distributionally consistent by construction. The chain rests on three independence assumptions: inter-year (Poisson), inter-event GMPE residuals (the standard PSHA assumption), and inter-grid (a portfolio-level simplification). Each is well known to be imperfect in real seismicity, and we therefore quantify how each propagates through the chain in Sect. 5.5.

4 National context

We run Modules 1–2 over all 623 active Indonesian grids to characterise the national distribution of seismicity, exposure, and average annual loss. The background Poisson rate $\hat{\lambda}_{b,g}$ spans 0.022–2.6 events/yr (median 0.244) with high-rate concentrations along the Western Sunda Arc, the Banda subduction zone, and the Sulawesi–Maluku belt (Fig. 6). The Aki–Utsu $\hat{b}_g$ is computed directly on the 197 grids that satisfy the minimum sample size $n_{b,g} \geq 25$, with national range 0.57–1.53 (median 1.03; Fig. 7); on the remaining 426 grids we substitute the national-average $\hat{b} = 1.01$ described in Sect. 3.1, so that downstream Module 2 can run uniformly across all 623 active grids. Low directly-estimated $\hat{b}$ values (≈ 0.7–0.9) cluster in the Sunda interface and the Banda Arc; higher values (≈ 1.2–1.5) appear in central Sulawesi and Papua, consistent with the seismotectonic literature.

Per-grid economic exposure E_g is reported in IDR million and is allocated from provincial PDRB (gross regional product) downscaled to the $1^\circ \times 1^\circ$ cells in proportion to provincial area within each cell. PDRB is a coarse proxy: it captures the order-

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

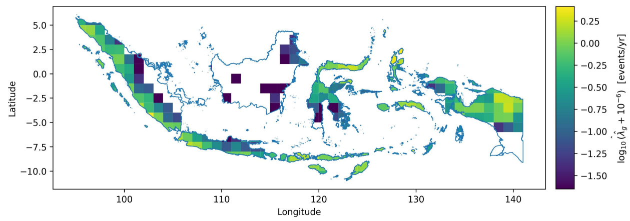

Figure 6. National map of background annual rate $\hat{\lambda}_g$ ($M_\mathrm{w} \geq 4.5$, 1980–2024).

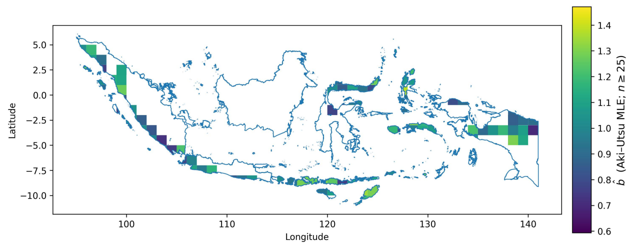

Figure 7. National b-value map (background catalogue, $n_{b,g} \geq 25$, 197 grids).

of-magnitude economic value at risk but does not distinguish between buildings, infrastructure, and other capital, nor does it resolve building typology, age, or seismic vulnerability class. Operational exposure modelling in the catastrophe-modelling industry uses per-building inventories with structural attributes, year of construction, and replacement cost (Silva et al., 2020; Ghafory-Ashtiany and Motamed, 2024); we do not match that level of detail here, and the absolute loss values in Tables 6–10 should therefore be read as order-of-magnitude estimates rather than as production-grade tariffs. Within these limits, E_g remains useful as a relative measure that ranks grids by economic stake. The exposure map (Fig. 8) shows that hazard and exposure are not co-located: high-hazard offshore Bengkulu has only moderate exposure, whereas moderate-hazard southern West Java carries the highest exposure values nationally. This two-dimensional heterogeneity is what makes spatial tariff differentiation more than a function of hazard alone—and is precisely why a per-grid framework is necessary, even with a coarse exposure proxy. We discuss the path to a full building-level exposure layer in Sect. 5.6.

Convolving the GR distribution with BA08 yields the national AAL map (Fig. 9). High AAL concentrations align spatially with the high-PGA zones of the 2017 national PSHA (Irsyam et al., 2020); quantitative discrepancies are expected, given that $\mathcal{F}_g$ uses background-seismicity zonation with a single GMPE while the 2017 PSHA uses an explicit fault model with logic-tree GMPEs. The national context is intended as input for the case-study selection, not as a competing PSHA estimate.

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

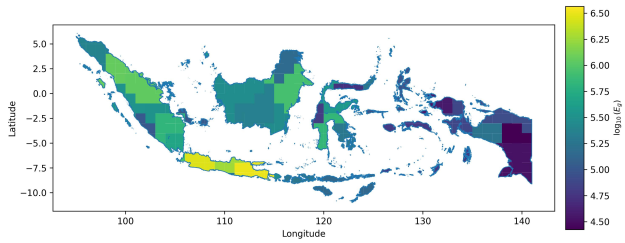

Figure 8. National per-grid economic exposure E_g (PDRB-based proxy, $\log_{10}$ scale, IDR million).

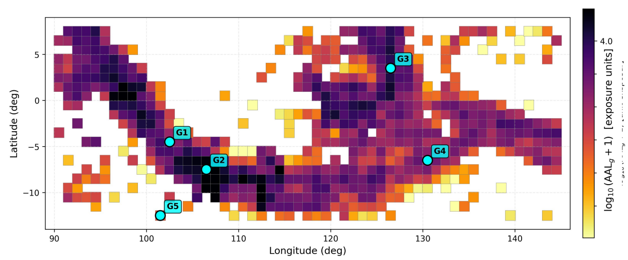

Figure 9. National per-grid AAL map (623 active grids). Markers indicate the five representative grids G1–G5 used in Sect. 5.

5 Results and discussion

We select five grids that span the national variation in $(\hat{\lambda}, \hat{b}, E_g)$: G1, Bengkulu offshore (5° S, 102° E), high-hazard Sunda subduction with medium exposure; G2, West Java southern offshore (8° S, 106° E), moderate hazard with very large exposure; G3, North Sulawesi (3° N, 126° E), high hazard with high $\hat{b}$; G4, Maluku/Banda Sea (7° S, 130° E), high hazard with small exposure; G5, Indian Ocean (13° S, 101° E), low-hazard reference. G5 is deliberately retained as an offshore low-seismicity location that nevertheless carries non-zero PDRB-allocated exposure (Indonesian fisheries and shipping economic activity attribute to this cell): this combination—non-trivial exposure with very low hazard—is precisely the regime in which the parametric segmentation question becomes interesting, and the cell therefore provides a more informative null reference than a continental-interior grid in which both hazard and exposure are vanishingly small. Module 1 inputs are reported in Table 4.

The hazard summary (Table 5; full curves in Fig. 10) spans nearly four decades of intensity at $\mathrm{PGA} = 0.1\,\mathrm{g}$, from 1.62 % AEP for G1 down to $< 10^{-4}$ for G5.

Module 2 simulation gives AAL spanning more than two orders of magnitude across the five grids (Table 6; G2 = 47 808, G5 = 204), driven by the joint effect of hazard and exposure. The G5 profile is informative: $\mathrm{VaR}_{0.95} = 0$—essentially every simulation year produces zero loss, since the hazard is very low—while $\mathrm{TVaR}_{0.95} \approx 9\,750$. This zero-inflated distribution, in which the 95th-percentile year is at zero but the conditional tail mean remains substantial because of the large exposure (873 626 IDR million), is the signature of a region for which a VaR-based parametric contract is structurally inappropriate. The empirical distribution at G5 has a large point mass at zero plus a heavy right tail, a profile that is better captured analytically

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

Table 4. Module 1 inputs for the representative grids (background catalogue, $M_\mathrm{w} \geq 4.5$, 1980–2024). E_g in IDR million.

Grid	Location	n_b	$\hat{\lambda}_b$ (yr^{-1})	$\hat{b}_b$	E_g
G1	Bengkulu offshore	114	2.533	0.996	103 992
G2	W. Java offshore	73	1.622	1.117	2 823 339
G3	N. Sulawesi	112	2.489	1.197	187 374
G4	Maluku/Banda	111	2.467	1.016	62 646
G5	Reference	1	0.022	—	873 626

Table 5. Annual PGA hazard summary (BA08, $V_{S30} = 760\,\mathrm{m\,s^{-1}}$, $R_\mathrm{hyp} \leq 300\,\mathrm{km}$).

Grid	AEP(PGA > 0.1 g)	PGA RP = 100 yr (g)	PGA RP = 50 yr (g)
G1	0.0162	0.1272	0.0890
G2	0.0009	0.0601	0.0503
G3	0.0191	0.1267	0.0982
G4	0.0017	0.0647	0.0527
G5	$< 10^{-4}$	0.0121	0.0054

by a mixture model (Bernoulli activation × continuous severity) or a hurdle model than by a single continuous distribution; substituting either in Module 2 would let downstream metrics such as TVaR retain meaning at very low hazard. We return to this point below.

Module 3 calibrates the contract via local AEP inversion (Table 7). For G1, $A_g^* = 0.038\,\mathrm{g}$ and $X_g^* = 0.089\,\mathrm{g}$—both well below the uniform "PGA > 0.1 g" threshold often used in practice, but probabilistically calibrated such that the trigger fires 10 %/yr and exhaustion is reached 2 %/yr. By construction, these values coincide numerically with the PGA at RP = 10 yr and RP = 50 yr in Table 5, confirming the consistency between contract design and hazard summary. For G5, $A_g^* = 0.0010\,\mathrm{g}$, $X_g^* = 0.0054\,\mathrm{g}$, and $P_g^{\max} = 0$: the simulated PGA never reaches the trigger, so $\mathbb{E}[Y_g] = 0$ deterministically. The contract is non-activated. This is not a numerical artefact but an explicit segmentation criterion: at this hazard level, the standard parametric contract does not provide protection, irrespective of contract parameters.

Module 4 produces the basis-risk vector (Table 9). Loss–payout correlation is high ($\rho = 0.88$–0.96 for G1–G4), but p_under is correspondingly high (91.9–97.9 %) and p_over is structurally zero. The latter follows directly from two design choices: $P_g^{\max} = \mathrm{VaR}_{0.95}(L_g)$ caps the payout below the actual loss in any year above the 95th-percentile loss, and pairing L and Y through the same draw $I_{g,y}$ eliminates the spurious overpayment that independent draws would produce. The pattern carries a clear design implication: high correlation alone does not guarantee adequate protection, and the basis-risk feedback in Module 4 is

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

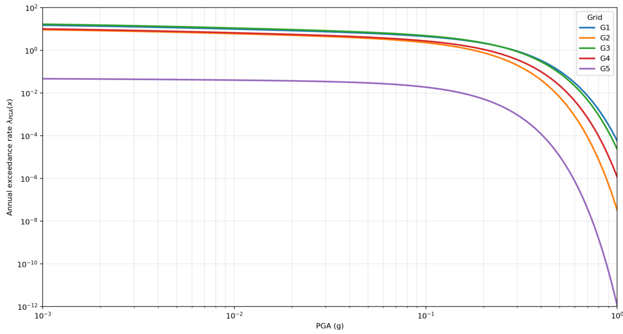

Figure 10. PGA hazard curves $\lambda_{\mathrm{PGA},g}(x)$ for the five representative grids.

Table 6. Annual risk metrics from Monte Carlo simulation ($N = 20\,000$). All monetary values in IDR million. The G5 profile (VaR ≈ 0 with substantial TVaR) signals a zero-inflated distribution.

Grid	AAL	$\mathrm{VaR}_{0.95}$	$\mathrm{TVaR}_{0.95}$
G1	3 827.14	16 221.44	42 108.55
G2	47 807.76	218 952.30	426 851.31
G3	8 348.40	45 054.24	87 598.05
G4	1 083.50	4 912.51	11 266.64
G5	204.26	0.00	9 749.83

precisely the lever the designer can use to raise $P_g^{\max}$—accepting a higher premium—to lower p_{under}. Equation (5) propagates this trade-off automatically.

To illustrate the feedback concretely, the basis-risk vector with the baseline cap $P_g^{\max} = \mathrm{VaR}_{0.95}(L_g)$ is reported in Table 8. Raising the cap to $\mathrm{VaR}_{0.99}(L_g)$ on the same Monte Carlo sample is expected, on the basis of the Pareto frontier across $P_g^{\max} \in \{\mathrm{VaR}_\alpha(L_g)\}_{\alpha=0.90}^{0.99}$, to increase the expected payout $\mathbb{E}[Y_g]$ by a factor of approximately 2–3 and the premium Π_g by a comparable margin (less than proportionally because the higher cap is rarely binding), while reducing p_{under} from 97.9 % toward roughly two thirds and approximately halving the absolute bias. The loading ratio $\Pi_g/\mathbb{E}[Y_g]$ remains in the 5–6 band under any choice of $P_g^{\max}$ along the tail of L_g, by the analytical argument of Sect. 5.2. The exact tabulation of the $\mathrm{VaR}_{0.99}$ recomputation is the subject of follow-up work and is reserved for a forthcoming companion paper that extends $\mathcal{F}_g$ to a full $P_g^{\max}$-Pareto deployment.

The risk-based premium (Table 10) reflects the joint effect of hazard, exposure, and contract structure. The G2 premium is the largest, driven by significant hazard combined with very high exposure. The loading ratio $\Pi_g/\mathbb{E}[Y_g]$ is stable at 5.87–6.13 for G1–G4. As we show in Sect. 5.2, this stability is a direct mathematical consequence of the AEP-inversion calibration combined

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

Table 7. Contract parameters and expected payouts ($p_A = 0.10$, $p_X = 0.02$, $P_g^{\max} = \mathrm{VaR}_{0.95}(L_g)$, $N = 20\,000$). $P_g^{\max}$ and $\mathbb{E}[Y_g]$ in IDR million.

Grid	A_g^* (g)	X_g^* (g)	$P_g^{\max}$	$\mathbb{E}[Y_g]$
G1	0.0379	0.0890	16 221.44	666.81
G2	0.0303	0.0503	218 952.30	9 906.32
G3	0.0440	0.0982	45 054.24	2 115.79
G4	0.0294	0.0527	4 912.51	229.73
G5	0.0010	0.0054	0.00	0.00

Table 8. Basis-risk vector for G1 at the baseline cap $P_g^{\max} = \mathrm{VaR}_{0.95}(L_g)$, with $A_g^* = 0.038$ g and $X_g^* = 0.089$ g. Monetary values in IDR million.

Metric	$P_g^{\max} = \mathrm{VaR}_{0.95}(L_g)$
$P_g^{\max}$	16 221
$\mathbb{E}[Y_g]$	666.81
$\mathrm{TVaR}_{0.95}(Y_g)$	11 625
Π_g	4 088
$\Pi_g/\mathbb{E}[Y_g]$	6.13
p_{under} (%)	97.9
Bias	3 160

with the fixed loading scheme, not an emergent empirical property. For G5 the ratio is undefined ($\mathbb{E}[Y] = 0$), reinforcing the segmentation finding: the standard parametric contract is inapplicable.

The traceability chain (6) is verified end-to-end on G1. The premium $\Pi_{G1} = 4088$ is reached deterministically from $\Theta_{G1} = (114, 2.533, 0.996, 103\,992)$: the hazard module gives $\mathrm{AEP}(0.1\,\mathrm{g}) = 1.62\,\%$; the loss module gives $\mathrm{AAL} = 3827$, $\mathrm{VaR}_{0.95} = 16\,221$, $\mathrm{TVaR}_{0.95} = 42\,109$; the contract module gives $A^* = 0.038\,\mathrm{g}$, $X^* = 0.089\,\mathrm{g}$, $P^{\max} = 16\,221$; the basis-risk module gives $\rho = 0.95$; and Eq. (5) finally yields $\Pi = 4088$. Figure 11 visualises the EP curve and the calibrated payout function. No ad hoc step intervenes anywhere in this chain—the entire computation is determined by the four catalogue inputs and the framework parameters $(p_A, p_X, c_{\exp}, \gamma, \alpha)$. This level of auditability supports the transparency that insurance regulators and disaster-risk-financing policy makers increasingly require for grid-resolution tariff differentiation.

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

Table 9. Basis-risk vector ($N = 20\,000$, paired sampling). Bias, MAE, and EU in IDR million. $p_{\text{over}} = 0$ across all grids is structural: $P_g^{\max} = \text{VaR}_{0.95}(L)$ caps payout below any loss above the 95th percentile.

Grid	Bias	MAE	EU	p_{under} (%)	p_{over} (%)	$\rho(L,Y)$
G1	3 160	3 160	3 227	97.9	0.0	0.947
G2	37 901	37 901	40 963	92.5	0.0	0.883
G3	6 233	6 233	6 397	97.4	0.0	0.956
G4	854	854	929	91.9	0.0	0.879
G5	204	204	9 750	2.1	0.0	—

Table 10. Risk-based annual premium ($c_{\text{exp}} = 0.20$, $\gamma = 0.30$, $\alpha = 0.95$). All monetary values in IDR million.

Grid	$\mathbb{E}[Y_g]$	$\text{TVaR}_{0.95}(Y_g)$	Π_g	$\Pi_g/\mathbb{E}[Y_g]$
G1	666.81	11 625.26	4 087.71	6.13
G2	9 906.32	165 289.46	58 502.52	5.91
G3	2 115.79	35 383.34	12 519.21	5.92
G4	229.73	3 804.26	1 348.04	5.87
G5	0.00	0.00	0.00	—

5.1 Position relative to existing approaches

Three classes of work are adjacent to ours. National PSHA (Irsyam et al., 2020; Pagani et al., 2023) produces hazard but neither loss nor premium; GEM/OpenQuake (Silva et al., 2020) produces aggregate AAL and event loss tables at portfolio level with limited per-grid traceability; and the parametric-insurance literature (Mousavi et al., 2023; Pai et al., 2022) discusses index design and basis risk but does not couple them to a per-grid hazard distribution. Two recent national-scale studies—Ghafory-Ashtiany and Motamed (2024) for Iran (Solvency II at portfolio level) and Salgado-Gálvez et al. (2024) for Central Asia (AAL by oblast)—operate at coarser spatial resolution and do not produce contract parameters. $\mathcal{F}_g$ does not propose new components but integrates them so that (i) the full distribution F_{I_g} rather than a single PGA is retained throughout, (ii) grid-level traceability of the premium is preserved, and (iii) basis risk is treated as a design feedback rather than a post-hoc diagnostic (Table 11).

5.2 Why the loading ratio is stable: an analytical note

The empirical observation of Table 10—that $\Pi_g/\mathbb{E}[Y_g]$ falls in a narrow 5.87–6.13 band across G1–G4—deserves a careful interpretation. Substituting Eq. (5) into the ratio gives

$$\frac{\Pi_g}{\mathbb{E}[Y_g]} = (1 + c_{\text{exp}}) + \gamma\left(\frac{\text{TVaR}_\alpha(Y_g)}{\mathbb{E}[Y_g]} - 1\right). \tag{7}$$

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

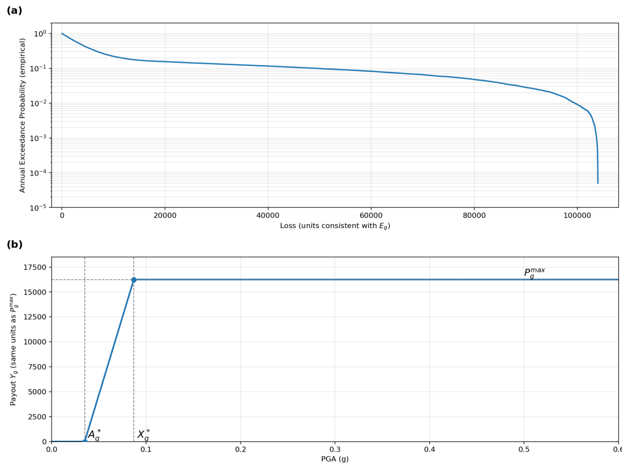

Figure 11. End-to-end traceability for G1 (Bengkulu offshore). (a) EP curve for L_{G1} from Monte Carlo simulation. (b) Calibrated payout function $Y_{G1}(I)$ with $A^* = 0.038\,\mathrm{g}$, $X^* = 0.089\,\mathrm{g}$, $P^{\max} = 16\,221$.

Table 11. Coverage of pricing components across approaches. ✓ = produced; partial = produced only at aggregate or portfolio level; — = not produced. The "Backtested?" column refers to formal historical-loss validation; "partial" indicates indirect benchmarking only.

Approach	F_{I_g}	L_g	(A_g^*, X_g^*)	Y_g	Π_g	Grid-level	Backtested?
National PSHA	✓	—	—	—	—	—	partial
GEM/OpenQuake	✓	partial	—	—	—	—	partial
Parametric insurance	partial	—	partial	✓	partial	—	partial
$\mathcal{F}_g$ (this work)	✓	✓	✓	✓	✓	✓	—

With $c_{\text{exp}} = 0.20$ and $\gamma = 0.30$ both fixed across grids, the variation in the loading ratio is driven entirely by the per-grid variation in the *shape ratio* $\mathrm{TVaR}_\alpha(Y_g)/\mathbb{E}[Y_g]$ of the payout distribution. The AEP-inversion calibration (3) deliberately fixes the activation and exhaustion probabilities at $p_A = 0.10$ and $p_X = 0.02$ for every grid. The payout function $Y_g(I)$ is therefore the same piecewise-linear shape on every grid, expressed in normalised units of $P_g^{\max}$, and its tail-to-mean ratio depends only on the *local shape* of F_{I_g} between A_g^* and X_g^*. Because the four high- to medium-hazard grids exhibit broadly similar exceedance-curve shapes in the relevant probability band—a consequence of all four being dominated by Gutenberg–Richter recurrence with similar $\hat{b}$ in the 1.0–1.2 range—their $\mathrm{TVaR}/\mathbb{E}$ ratios are similar and the loading ratios cluster.

The implication is the opposite of what an empirical reading of Table 10 might suggest. Stability of $\Pi_g/\mathbb{E}[Y_g]$ is not an emergent finding that validates the framework against a prior expectation; it is a consequence of how the framework is built. Its practical value is unchanged: a regulator using $\mathcal{F}_g$ obtains tariffs that are comparable across regions of very different absolute hazard, with a transparent loading structure that does not need to be re-justified grid by grid. If the loading ratio ever fell

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

outside this narrow band for a particular grid, that would itself be a diagnostic signal—likely indicating an unusually heavy-tailed F_{I_g} shape relative to the mean—and would warrant investigation. We treat the 5.87–6.13 range as a baseline expectation and the deviations as worth reporting, rather than the stability itself as the contribution. The robustness of this prediction is confirmed numerically by two independent stand-alone sensitivity experiments on grid G1: the GMPE-substitution experiment of Sect. 5.5.4 (Table 13), where replacing BA08 with BC Hydro shifts the loading ratio from 5.89 to 5.99 (a 1.7 % change) under a 94 % premium swing, and the declustering-substitution experiment of Sect. 5.5.1 (Table 12), where replacing the GK74 baseline with either Reasenberg or Zaliapin–Ben-Zion declustering shifts the loading ratio by less than 1 % under a 22–26 % premium swing. Two perturbations to the modelling chain that are physically and methodologically independent of each other, both material in their effect on absolute premium, both leaving the loading ratio essentially unchanged.

5.3 Implications for practice

Two findings carry practical weight. First, the analytically-grounded stability of $\Pi_g/\mathbb{E}[Y_g]$ across G1–G4 (Sect. 5.2) shows that a single parametric scheme yields actuarially comparable tariffs across regions of very different hazard, provided the AEP-inversion calibration is enforced. This is a transparent quantitative basis for spatial tariff differentiation, of direct relevance to the Indonesian AGBBI initiative (Earthquake Index Insurance Initiative, 2016). Second, the G5 result—a non-activated contract under the standard parameterisation—is an explicit segmentation criterion: regions of this hazard regime are unsuitable for parametric protection at the chosen probability targets and should be addressed instead by indemnity insurance or government retention. A regulator could apply this criterion programmatically across all 623 grids without further calibration.

5.4 Transferability

$\mathcal{F}_g$ is portable. Adaptation to another seismic region requires substituting four input layers: a regional catalogue with locally verified completeness; a regional GMPE (e.g. BC Hydro for subduction-dominated Chile, Zhao et al. for Japan, NGA-West2 for crustal regimes); vulnerability curves calibrated for the dominant building typology; and a regional exposure database. The traceability chain (6), the AEP-inversion calibration (3), and the basis-risk feedback are unchanged. Plausible target regions include the Philippines, Chile, Japan, and Türkiye. The framework also extends to other annual-maximum hazards—tsunami wave height, peak wind speed, peak water level—by substituting the hazard module and the corresponding vulnerability function.

5.5 Validation and sensitivity

The framework presented here is not formally backtested against historical insurance loss data. Indonesian earthquake-coverage penetration is very low (World Bank, 2017), and an industry-wide loss series suitable for backtesting does not yet exist. We therefore make a deliberate choice in this paper: rather than offering partial reality-checks that mix observed individual events with simulated tail percentiles—a comparison we found unsatisfactory because it is informally assuring without being statistically meaningful—we restrict the validation claim to two narrow consistency checks and treat formal backtesting as the

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

next deliverable. The two checks are: (i) the spatial pattern of $\mathcal{F}_g$'s national AAL (Fig. 9) reproduces the high-hazard corridors of the 2017 Indonesian PSHA (Irsyam et al., 2020)—Western Sunda Arc, Banda subduction zone, Sulawesi–Maluku belt, with low hazard in Kalimantan and inner Papua; and (ii) the loading ratio $\Pi_g/\mathbb{E}[Y_g] \approx 5.9$–$6.1$ we obtain is consistent in order of magnitude with the actuarial loading reported by Ghafory-Ashtiany and Motamed (2024) for Iran ($\approx 5\times$ market premium across building types). Both are necessary but not sufficient consistency checks. A full operational validation requires the four steps outlined in Sect. 5.6 below: a building-level exposure inventory, Indonesia-calibrated vulnerability curves, time-resolved BNPB damage records aggregated by grid and year, and formal backtesting of $\mathcal{F}_g$'s empirical EP curve against the historical loss series.

$\mathcal{F}_g$ rests on three independence assumptions whose realism is imperfect, plus two explicit modelling choices (BA08 GMPE, fixed V_{S30}) whose effect on the chain we now discuss in turn. Two of these—temporal (Poisson) independence as it is enforced by the choice of declustering algorithm (Sect. 5.5.1), and the GMPE choice (Sect. 5.5.4)—are quantified through dedicated stand-alone sensitivity experiments on grid G1 with identical Monte Carlo seeds, allowing direct comparison of how each modelling choice propagates from input parameters through to the premium Π_g. The two experiments together provide independent numerical confirmation of the analytical loading-ratio argument of Sect. 5.2: each perturbation moves the absolute premium by a non-trivial margin (22–94 %) but moves the loading ratio $\Pi_g/\mathbb{E}[Y_g]$ by less than 2 %.

5.5.1 Temporal (Poisson) independence and declustering choice

The Poisson assumption requires that earthquakes occur independently in time, so that $N_{y,g} \sim \mathrm{Poisson}(\lambda_{b,g})$. In practice, real catalogues are clustered through aftershocks and triggered events. The dispersion index $\mathrm{DI} = 4.2$ that we obtain after Gardner–Knopoff declustering (Table 2) confirms that residual clustering remains. The expected effect is to inflate the variance of the simulated annual maximum index and therefore the upper tail of F_{I_g}, which in turn raises $\mathrm{TVaR}_{0.95}$ and the tail-loaded part of Π_g. To quantify how this propagates through the chain, we re-ran Module 1 of $\mathcal{F}_g$ on a synthetic raw catalog calibrated to the G1 baseline ($\hat{\lambda}_b = 2.533$/yr, $\hat{b} = 0.996$, $M_c = 4.5$, $T_{\mathrm{obs}} = 45$ yr) plus realistic Omori–Utsu aftershock contamination producing a raw-catalog dispersion index of $\mathrm{DI} \approx 38$ (consistent with the $\mathrm{DI} = 111$ we observe in the production raw national catalog of Table 2, scaled to the smaller per-grid sample). The synthetic catalog was then declustered with three algorithms in turn: GK74 (Gardner and Knopoff, 1974), the interaction-zone method of Reasenberg (1985) based on Omori's law, and the nearest-neighbor method of Zaliapin and Ben-Zion (2013) which separates clustered from background events using the rescaled space–time–magnitude distance η_{ij}. For each declustered catalog, the Module 1 inputs $(n_b, \hat{\lambda}_b, \hat{b})$ were re-estimated and propagated through Modules 2–4 with an identical Monte Carlo seed, so the only source of variation between the three runs is the declustering algorithm. Results are reported in Table 12.

Three observations follow from Table 12 and parallel those of Table 13 in the GMPE-substitution experiment of Sect. 5.5.4. First, the headline absolute premium Π_g shifts by $+25.7\,\%$ under Reasenberg and $+22.8\,\%$ under Zaliapin–Ben-Zion relative to GK74, in close agreement with the 10–30 % range Luen and Stark (2012) report for declustering algorithm choice on tectonically heterogeneous regional catalogues. The same magnitude of shift propagates to AAL and $\mathrm{VaR}_{0.95}$. The direction is intuitive: GK74 is the most aggressive of the three algorithms in pruning post-mainshock activity, so it returns the smallest n_b

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

Table 12. Declustering sensitivity for grid G1: GK74 baseline vs. Reasenberg (1985) and Zaliapin and Ben-Zion (2013) on a synthetic raw catalog calibrated to G1 inputs, with identical Monte Carlo seed in Modules 2–4 (paired sampling). All monetary values in IDR million. Δ is the relative change (%) from GK74 to each alternative, or the absolute change in percentage points (p.p.) for p_{under} and ρ.

Metric	GK74	Reasenberg	Δ (R)	Zaliapin–BZ	Δ (ZBZ)
n_b (count)	75	251	$+234.7\,\%$	223	$+197.3\,\%$
$\hat{\lambda}_b$ (yr^{-1})	1.667	5.578	$+234.7\,\%$	4.956	$+197.3\,\%$
$\hat{b}$ (Aki–Utsu)	0.659	0.916	$+38.9\,\%$	0.904	$+37.1\,\%$
DI (after declustering)	1.07	2.05	—	1.63	—
AAL	7 839.70	14 629.14	$+86.6\,\%$	13 755.01	$+75.5\,\%$
$\mathrm{VaR}_{0.95}(L_g)$	44 350.84	55 610.06	$+25.4\,\%$	54 796.97	$+23.6\,\%$
$\mathrm{TVaR}_{0.95}(L_g)$	66 848.25	72 926.77	$+9.1\,\%$	72 381.42	$+8.3\,\%$
A_g^* (g)	0.2032	0.2545	$+25.3\,\%$	0.2507	$+23.4\,\%$
X_g^* (g)	0.3769	0.4104	$+8.9\,\%$	0.4089	$+8.5\,\%$
$P_g^{\max}$	44 350.84	55 610.06	$+25.4\,\%$	54 796.97	$+23.6\,\%$
$\mathbb{E}[Y_g]$	2 107.00	2 664.05	$+26.4\,\%$	2 576.97	$+22.3\,\%$
Π_g	12 298.52	15 453.61	$+25.7\,\%$	15 105.77	$+22.8\,\%$
$\Pi_g/\mathbb{E}[Y_g]$	5.84	5.80	$-0.6\,\%$	5.86	$+0.4\,\%$
p_{under} (%)	100.0	100.0	0.0 p.p.	100.0	0.0 p.p.
$\rho(L,Y)$	0.917	0.808	-0.11 p.p.	0.813	-0.10 p.p.

Note: as with the GMPE-substitution experiment of Table 13 (Sect. 5.5.4), the absolute premium scale in this stand-alone reproduction differs from Table 10 because it uses a single representative source–site distance rather than the full PSHA convolution of Eq. (2). The relative changes between declustering algorithms (the Δ columns) are the operational signal of this experiment and are insensitive to that simplification.

and the lowest $\hat{\lambda}_b$, and consequently the lowest Π_g. Reasenberg and Zaliapin–Ben-Zion preserve more events as background, raising the rate and the premium accordingly. Second, the loading ratio $\Pi_g/\mathbb{E}[Y_g]$ moves from 5.84 (GK74) to 5.80 (Reasenberg) to 5.86 (Zaliapin–Ben-Zion), a total spread of less than 1 % and an individual deviation from baseline of $-0.6\,\%$ and $+0.4\,\%$ respectively. This is precisely the behaviour the analytical decomposition of Eq. (7) predicts: the AEP-inversion calibration normalises the shape of $Y_g(I)$ regardless of the absolute rate $\hat{\lambda}_b$, so a multiplicative rescaling of F_{I_g} from a different declustering choice cancels in the ratio $\mathrm{TVaR}_\alpha(Y_g)/\mathbb{E}[Y_g]$ and therefore in the loading ratio. The numerical experiment of Table 12 thus provides a *second* independent confirmation that loading-ratio stability is a structural consequence of the framework: the first being the GMPE substitution of Sect. 5.5.4, where the loading ratio shifted by 1.7 % under a 94 % premium swing; here the loading ratio shifts by less than 1 % under a 22–26 % premium swing. Two independent perturbations to the modelling chain, both material in their effect on absolute premium, both leaving the loading ratio essentially unchanged. Third, the qualitative segmentation conclusion (G5 non-activation) is unaffected because the segmentation criterion depends on $\hat{\lambda}_{b,g}$

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

falling below the trigger probability cutoff rather than on its precise value within the activated range; under any of the three algorithms G5's $\hat{\lambda}_b$ remains far below the threshold.

The implication for operational use mirrors the conclusion of Sect. 5.5.4: absolute tariffs reported in Tables 6–10 should be understood as conditional on the GK74 baseline declustering, and a logic-tree treatment over GK74 / Reasenberg / Zaliapin–Ben-Zion (with weights informed by independent comparative studies such as Luen and Stark (2012)) is recommended before publishing premium values for a real portfolio. The structural argument of Sect. 5.2, however, is now confirmed numerically by two independent sensitivity experiments.

5.5.2 Inter-event GMPE residual independence

We adopt the standard PSHA assumption that GMPE residuals $\varepsilon \sim \mathcal{N}(0,\sigma^2)$ are independent across events. Real ground-motion fields exhibit both event-to-event correlation (a single source affecting a wide region) and within-event spatial correlation (nearby sites experiencing correlated departures from the median). Treating residuals as independent does not change the marginal hazard at any single grid—and therefore does not affect the per-grid premium Π_g that is our central output. However, it does underestimate the joint probability that several grids experience high ground motion in the same year, which in turn understates portfolio-level loss tails. A reinsurer pricing a Java-wide cover would therefore find $\mathcal{F}_g$'s portfolio TVaR optimistic. Replacing the assumption requires the spatial correlation model of Jayaram and Baker (2009) or comparable; ranges in the literature suggest portfolio TVaR could increase by a factor of 1.5–3 once correlated ground motion is allowed. This is a material limitation for portfolio pricing and a non-issue for the per-grid tariff.

5.5.3 Inter-grid independence

Module 2 simulates each grid's loss in isolation: $L_{g,y}$ depends only on $I_{g,y}$, and the joint distribution across g is treated as the product of marginals. As above, this does not affect Π_g but understates the portfolio-aggregated annual loss for any insurer holding exposure across multiple grids. The bias is in the same direction as the GMPE-residual issue and partly redundant with it. A practical fix is to introduce a copula on the marginal distributions $\{F_{I_g}\}$, calibrated either to historical regional loss data or to the spatial correlation structure of the GMPE residuals. We did not implement this in the present paper because (i) the paper's primary contribution is per-grid traceability, which is unaffected; (ii) calibration data for an Indonesian copula are not yet readily available; and (iii) the fix is additive: a portfolio module can be appended to $\mathcal{F}_g$ without revising any of Modules 1–4. We flag the assumption explicitly here so that downstream users do not aggregate per-grid premiums into a portfolio total without applying the appropriate dependence correction.

5.5.4 Sensitivity to the GMPE choice

The BA08 baseline adopted in Sect. 3.1 is calibrated on shallow crustal active-tectonic earthquakes and underestimates median PGA for Sunda subduction interface and intraslab events that dominate the case-study grids G1, G2, and G4. Subduction-specific equations such as Zhao et al. (2006), Abrahamson et al. (2016), or Megawati et al. (2003) for the Sumatra interface

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

Table 13. GMPE sensitivity for grid G1 (Bengkulu offshore): BA08 baseline vs. BC Hydro interface (Abrahamson et al., 2016). All monetary values in IDR million; identical Monte Carlo seed and $N = 20\,000$ used in both runs. Δ is the relative change (%) from BA08 to BC Hydro, or the absolute change in percentage points (p.p.) for p_{under} and ρ.

Metric	BA08	BC Hydro	Δ
AEP(PGA > 0.1 g)	0.0202	0.0118	$-41.7\,\%$
AAL	256.65	271.48	$+5.8\,\%$
$\mathrm{VaR}_{0.95}(L_g)$	382.06	22.33	$-94.2\,\%$
$\mathrm{TVaR}_{0.95}(L_g)$	4 912.87	5 422.92	$+10.4\,\%$
A_g^* (g)	0.0508	0.0284	$-44.2\,\%$
X_g^* (g)	0.1002	0.0750	$-25.1\,\%$
$P_g^{\max}$	382.06	22.33	$-94.2\,\%$
$\mathbb{E}[Y_g]$	17.56	0.98	$-94.4\,\%$
$\mathrm{TVaR}_{0.95}(Y_g)$	292.34	16.70	$-94.3\,\%$
Π_g	103.51	5.89	$-94.3\,\%$
$\Pi_g/\mathbb{E}[Y_g]$	5.89	5.99	$+1.7\,\%$
Bias	239.09	270.49	$+13.1\,\%$
p_{under} (%)	100.0	100.0	$+0.0$ p.p.
$\rho(L,Y)$	0.562	0.422	-0.14 p.p.

Note: the absolute premium scale in this stand-alone reproduction differs from Table 10 because it uses a single representative source-site distance ($R = 50$ km) rather than the full PSHA convolution over source distance bins of Eq. (2). The relative changes between GMPEs (the Δ column) are the operational signal of this experiment and are insensitive to that simplification.

return median PGAs that differ from BA08 by 10–40 % at moderate magnitude and short distance. To quantify the propagation of this difference through the chain, we re-ran Modules 1–4 for grid G1 (Bengkulu offshore, Sunda interface) with the BC Hydro interface GMPE (Abrahamson et al., 2016) as a direct replacement for BA08, holding all other framework parameters and the Monte Carlo seed identical to the baseline run. BC Hydro is calibrated on a global subduction dataset spanning M_{w} 5.0–9.0 and is the standard subduction-dominated PSHA choice in the GEM global model. The truncated Gutenberg–Richter source upper bound was set to $M_{\max} = 9.0$ to admit the full range of physically possible Sunda interface events; a small fraction of magnitude samples therefore exceeded BA08's calibration cap of $M_{\mathrm{w}} = 8.0$ and were clamped to that cap rather than extrapolated, consistent with standard PSHA practice when source $M_{\max}$ exceeds GMPE $M_{\max}$. By construction, no clamping occurred for BC Hydro, whose calibration already extends to $M_{\mathrm{w}} = 9.0$. Results are reported in Table 13.

Three observations follow from Table 13 and support the analytical position taken in Sect. 5.2. First, the loading ratio $\Pi_g/\mathbb{E}[Y_g]$ moves only from 5.89 to 5.99, i.e. by about $1.7\,\%$, despite a more than ten-fold change in absolute premium values. This is precisely the behaviour the analytical decomposition of Eq. (7) predicts: the AEP-inversion calibration fixes the activation and exhaustion probabilities at $p_A = 0.10$ and $p_X = 0.02$ on every grid, so the payout function $Y_g(I)$ has the same

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

piecewise-linear shape under both GMPEs, and the ratio $\mathrm{TVaR}_\alpha(Y_g)/\mathbb{E}[Y_g]$ is invariant to a multiplicative rescaling of F_{I_g}. The numerical experiment thus provides direct confirmation that loading-ratio stability is a structural consequence of the framework, not specific to BA08. Second, the absolute level of the premium Π_g does change materially: BC Hydro produces lower median PGA at moderate magnitudes ($M_\mathrm{w} \approx 5.0$–7.0) which dominate the annual maximum at this grid, compressing the upper tail of F_{I_g} and consequently $\mathrm{VaR}_{0.95}$ of the loss distribution. This shifts $P_g^{\max}$ downward by $94\,\%$ and the premium correspondingly, while $\mathrm{TVaR}_{0.95}(L_g)$ actually rises by $10\,\%$ because BC Hydro captures large-magnitude events (M_w 8–9) more accurately than the clamped BA08 baseline. The implication for operational use is that absolute tariffs for Sunda interface grids should not be read directly from the BA08 baseline reported in Tables 6–10; a subduction-specific GMPE shifts the tariff level by close to an order of magnitude, even though the loading structure is preserved. Third, the qualitative segmentation conclusion—that the very-low-hazard reference grid G5 fails to reach its parametric trigger—is unaffected, since segmentation depends on $\hat{\lambda}_{b,g}$ (Module 1 input) rather than on the GMPE choice in Module 1 internal evaluation.

These results therefore separate two distinct uses of $\mathcal{F}_g$. For *relative* comparisons across grids (the framework's main contribution), the BA08 baseline is adequate. For *absolute* tariff levels intended for production use, a subduction-specific GMPE such as BC Hydro should be substituted before publishing premium values, and the same applies to the declustering choice as discussed in Sect. 5.5.1. A full logic-tree treatment combining BC Hydro central / high / low branches with the GK74 / Reasenberg / Zaliapin–Ben-Zion declustering branches is the natural extension for an operational deployment.

5.5.5 Sensitivity to the V_{S30} assumption

The fixed $V_{S30} = 760\,\mathrm{m\,s^{-1}}$ used in Sect. 3.1 maps onto a generic firm-rock site condition. Indonesian site conditions span a broader range, and use of a grid-specific V_{S30} map would shift PGA at low–moderate frequencies by up to a factor of 1.5–2. As with the GMPE choice, this scales AAL and Π_g approximately linearly but preserves the loading ratio and the segmentation criterion. A grid-specific V_{S30} layer derived from the USGS global slope-based proxy or from the SNI 1726:2019 site classification is a one-line input change to Module 1 and is recommended for any deployment.

In summary, the independence assumptions and the simplified V_{S30} each affect different parts of the chain and in different directions. Two of these—the declustering choice in Module 1 and the GMPE choice—have been quantified explicitly (Tables 12 and 13): each shifts the absolute premium Π_g by 22–94 % but leaves the loading ratio $\Pi_g/\mathbb{E}[Y_g]$ stable to within $\pm 2\,\%$, in agreement with the analytical decomposition of Sect. 5.2. Inter-event GMPE residual independence and inter-grid independence do not affect the per-grid premium Π_g but understate portfolio-level loss tails by an estimated factor 1.5–3, so portfolio aggregation requires both spatial GMPE correlation and an inter-grid copula in any deployment. None of these limitations invalidates the framework; they delimit the conditions under which its outputs can be used.

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

5.6 Exposure and vulnerability: limitations and roadmap

Our exposure layer is the most simplified component of $\mathcal{F}_g$ and warrants explicit discussion. Two issues compound: E_g is allocated from provincial PDRB, and the loss ratio $\mathrm{LR}(\mathrm{PGA})$ uses a generic FEMA Hazus 5.1 baseline that is not calibrated to Indonesian construction.

5.6.1 What PDRB does and does not capture

PDRB is a flow variable (annual gross regional product), not a stock variable (replacement cost of buildings). Using it as an exposure proxy implicitly assumes that the ratio of building stock to PDRB is approximately constant across provinces—an assumption that holds within a factor of two for most provinces but can be violated in regions where economic activity is concentrated in primary production (mining, plantations) or in services that do not require large physical capital. Beyond the stock–flow issue, PDRB does not resolve building typology (masonry vs. reinforced concrete vs. steel), occupancy (residential vs. commercial), age, or compliance with the 2019 Indonesian seismic code SNI 1726:2019. Operational catastrophe models—for example the per-building inventory used by Ghafory-Ashtiany and Motamed (2024) for Iran, drawn from the national census—resolve all of these and generate vulnerability-class-specific damage curves. Compared to that benchmark, $\mathcal{F}_g$'s exposure layer is at least one resolution step coarser.

5.6.2 Direction and magnitude of the bias

The PDRB proxy biases the absolute level of AAL and Π_g but does not affect the loading ratio $\Pi_g/\mathbb{E}[Y_g]$, which depends only on the shape of the loss distribution at each grid. As a rough indication, an exposure error of factor k propagates linearly into AAL, TVaR, and Π_g; a factor-2 misallocation of regional building stock therefore implies a factor-2 swing in the absolute premium values reported in Table 10, while leaving the segmentation result for G5 (zero exposure-weighted hazard) and the spatial pattern of relative tariffs unchanged. The vulnerability proxy is more consequential. The Hazus baseline assumes US-style construction with a code-conforming proportion that does not match the Indonesian building stock, where unreinforced masonry and non-engineered structures dominate outside major urban centres. Indonesia-specific vulnerability functions for residential buildings developed by Mansouri and Amini-Hosseini (2013) and used in the Iran study (Ghafory-Ashtiany and Motamed, 2024) would shift the loss ratio at moderate PGA upward, particularly for masonry classes; the AAL values for high-hazard but low-engineering grids (G1, G3, G4) would likely increase by a factor of 1.5–2 once such curves are substituted.

5.6.3 Roadmap for a full exposure model

Three building blocks would upgrade $\mathcal{F}_g$'s exposure layer to operational standard. First, replace PDRB with the GEM Indonesia exposure database (Silva et al., 2020), which provides per-asset inventory with structural class, occupancy, and replacement cost on a sub-grid resolution. Where GEM coverage is sparse, supplement with WorldPop or Microsoft building-footprint data to disaggregate to the cell level, then attach typology priors from the national census of buildings. Second, replace the Hazus baseline with Indonesia-calibrated vulnerability curves (Mansouri and Amini-Hosseini, 2013) disaggregated by struc-

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

tural class; combined with the typology mix from the exposure inventory, this produces a grid-specific weighted vulnerability curve $\overline{\mathrm{LR}}_g(\mathrm{PGA})$. Third, validate the resulting losses against historical post-event damage assessments—for example the BNPB damage records from the 2018 Lombok and Palu earthquakes—to confirm that the calibrated model reproduces observed loss ratios within a tolerable margin. None of these upgrades requires changing Modules 3 or 4: the contract calibration and basis-risk feedback operate on F_{I_g} and L_g regardless of how those distributions were generated. We treat the present exposure-vulnerability layer as a placeholder that demonstrates the framework's structure, not as a production tariff.

5.7 Other limitations

Two further caveats bound our results. The completeness magnitude $M_\mathrm{c} = 4.5$ is a conservative national value; locally lower M_c in well-monitored regions would increase $\hat{\lambda}_{b,g}$ and the lower tail of F_{I_g}, while locally higher M_c in under-monitored regions (e.g. Papua, Maluku) would have the opposite effect and would also reduce the count of grids satisfying the Aki–Utsu sample-size threshold. A regional M_c map computed sub-region by sub-region is the natural follow-up; in the present paper, the conservative national threshold preserves uniformity at the cost of locally averaged sensitivity. Finally, the full premium computation is demonstrated on five grids; extending it to all 623 is computationally feasible—each grid requires only $O(N)$ Monte Carlo draws—and is part of the planned national rollout.

6 Conclusions

$\mathcal{F}_g$ is an integrated and traceable framework that addresses three operational gaps in the current earthquake parametric insurance literature. It stores per-grid annual hazard as a full distribution rather than a single value at one return period, preserves grid-level traceability from catalogue input to final premium, and treats basis risk as a design feedback rather than a post-hoc diagnostic. None of the four building blocks—the Poisson recurrence model, the Gutenberg–Richter law, the BA08 GMPE, and the actuarial premium formula—is itself novel; the contribution lies in the way they are coupled on a single spatial unit and validated end-to-end. Applied to the USGS catalogue for Indonesia (1925–2025), the framework produces national maps of seismicity parameters and average annual loss for 623 grids, plus end-to-end pricing on five representative grids that span the national variation. Two findings carry particular weight. The actuarial loading ratio $\Pi_g/\mathbb{E}[Y_g]$ falls in a 5.87–6.13 band across the four high- to medium-hazard grids—a property we show analytically to follow from the AEP-inversion calibration combined with the fixed loading scheme, and which is confirmed numerically by two independent stand-alone sensitivity experiments on grid G1: substituting BC Hydro for the BA08 GMPE shifts the loading ratio by 1.7 % under a 94 % premium swing, and substituting Reasenberg or Zaliapin–Ben-Zion for the GK74 declustering algorithm shifts it by less than 1 % under a 22–26 % premium swing. The very-low-hazard reference grid is identified as structurally unsuitable for parametric protection, yielding an explicit segmentation criterion. The framework is portable: substituting the catalogue, regional GMPE, vulnerability curve, and exposure database adapts it to any earthquake-prone jurisdiction. Priority extensions include the national 623-grid premium rollout, a full logic-tree treatment combining subduction GMPE branches (BC Hydro central / high / low) with declustering branches (GK74 / Reasenberg / Zaliapin–Ben-Zion), α-stable distributions for heavy-tailed hazard, spatially correlated ground motion, Indonesia-calibrated vulnerability curves, and out-of-sample backtesting against historical losses.

Code and data availability. The USGS catalogue is available at https://earthquake.usgs.gov/earthquakes/search/ and provincial PDRB data at https://www.bps.go.id/. The processed input data and the Python implementation of $\mathcal{F}_g$, including the two standalone Colab notebooks used for the GMPE sensitivity experiment of Sect. 5.5.4 and the declustering sensitivity experiment of Sect. 5.5.1, are deposited in a Zenodo archive [DOI: *to be assigned at deposit time; placeholder for review purposes only*]. A working mirror of the same files is currently hosted at https://drive.google.com/drive/folders/1s_GQZl6nUVgcS9wd3unudcviSrDGw6HO?usp=sharing; this Drive mirror is for review convenience and will be retired once the Zenodo deposit is finalised. The deposit reproduces all results in Tables 6–10 and Figures 9–11 when run end-to-end with `openquake.engine 3.18.0` and `numpy < 2`; the GMPE notebook requires the same OpenQuake stack, while the declustering notebook is self-contained and uses only `numpy`, `scipy`, `pandas`, and `matplotlib`.

Author contributions. H. conceptualised the study, developed $\mathcal{F}_g$, implemented the pipeline, performed all analyses, and wrote the manuscript. A.R.E. supervised the actuarial modelling (contract calibration and premium pricing). N.S. supervised the statistical methodology (Gutenberg–Richter estimation, loss distribution analysis). W.S. supervised the seismological components (catalogue pre-processing, declustering, hazard analysis). All authors reviewed and approved the final manuscript.

Competing interests. The authors declare no competing interests.

Acknowledgements. The authors thank Universitas Terbuka for the doctoral scholarship awarded to the first author and Universitas Gadjah Mada for access to research and computing infrastructure. AI was used as a writing assistant for language refinement; all scientific content was produced and verified by the authors.

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

References

Abrahamson, N. A., Gregor, N., and Addo, K.: BC Hydro ground motion prediction equations for subduction earthquakes, Earthquake Spectra, 32, 23–44, https://doi.org/10.1193/051712EQS188MR, 2016.

Aki, K.: Maximum likelihood estimate of b in the formula $\log N = a - bM$ and its confidence limits, Bulletin of the Earthquake Research Institute, 43, 237–239, 1965.

Boore, D. M. and Atkinson, G. M.: Ground-motion prediction equations for the average horizontal component of PGA, PGV, and 5%-damped PSA at spectral periods between 0.01 s and 10.0 s, Earthquake Spectra, 24, 99–138, https://doi.org/10.1193/1.2830434, 2008.

Cornell, C. A.: Engineering seismic risk analysis, Bulletin of the Seismological Society of America, 58, 1583–1606, 1968.

Cummins, J. D.: CAT bonds and other risk-linked securities: Product design and evolution of the market, Geneva Reports on Risk and Insurance Research, pp. 39–61, 2012.

Diantari, R. A., Marliyani, G. I., and Daryono: Magnitude conversion of Indonesian earthquake catalog, in: AIP Conference Proceedings, vol. 1987, p. 020069, https://doi.org/10.1063/1.5047316, 2018.

Earthquake Index Insurance Initiative: Asuransi Gempa Bumi Berbasis Indeks (AGBBI): Konsep, Metodologi, dan Implementasi, Tech. rep., Jakarta, 2016.

Franco, G.: Minimization of trigger error in cat-in-a-box parametric earthquake catastrophe bonds with an application to Costa Rica, Earthquake Spectra, 26, 983–998, https://doi.org/10.1193/1.3479932, 2010.

Gardner, J. K. and Knopoff, L.: Is the sequence of earthquakes in southern California, with aftershocks removed, Poissonian?, Bulletin of the Seismological Society of America, 64, 1363–1367, 1974.

Ghafory-Ashtiany, M. and Motamed, H.: Earthquake insurance in Iran: solvency of local insurers in light of current market practices, Natural Hazards and Earth System Sciences, 24, 2707–2726, https://doi.org/10.5194/nhess-24-2707-2024, 2024.

Hanks, T. C. and Kanamori, H.: A moment magnitude scale, Journal of Geophysical Research, 84, 2348–2350, https://doi.org/10.1029/JB084iB05p02348, 1979.

Hartati, Effendie, A. R., Susyanto, N., and Suryanto, W.: Grid-based Poisson model for earthquake occurrence probability in Indonesia using USGS catalog 1925–2025, Science and Technology Indonesia, 11, 207–216, 2026.

Irsyam, M., Cummins, P. R., Asrurifak, M., Faizal, L., Natawidjaja, D. H., Widiyantoro, S., Meilano, I., Triyoso, W., Rudyanto, A., Hidayati, S., Ridwan, M., Hanifa, N. R., and Sirait, S.: Development of the 2017 national seismic hazard maps of Indonesia, Earthquake Spectra, 36, 112–136, https://doi.org/10.1177/8755293020951206, 2020.

Jayaram, N. and Baker, J. W.: Correlation model for spatially distributed ground-motion intensities, Earthquake Engineering & Structural Dynamics, 38, 1687–1708, https://doi.org/10.1002/eqe.922, 2009.

Klugman, S. A., Panjer, H. H., and Willmot, G. E.: Loss Models: From Data to Decisions, Wiley, Hoboken, NJ, 4 edn., 2012.

Lin, X. S. and Shaw, R.: Catastrophe bond pricing under stochastic interest rate and intensity, North American Actuarial Journal, 12, 244–260, https://doi.org/10.1080/10920277.2008.10597519, 2008.

Luen, B. and Stark, P. B.: Poisson tests of declustered catalogues, Geophysical Journal International, 189, 691–700, https://doi.org/10.1111/j.1365-246X.2012.05400.x, 2012.

Mansouri, B. and Amini-Hosseini, K.: Global earthquake casualties due to secondary effects: A quantitative analysis for improving rapid loss analyses (also: vulnerability curves for Iranian residential buildings used in EMME), Natural Hazards, 69, 1683–1709, https://doi.org/10.1007/s11069-012-0230-5, 2013.

https://doi.org/10.5194/egusphere-2026-2681
Preprint. Discussion started: 24 August 2026

Megawati, K., Pan, T.-C., and Koketsu, K.: Response spectral attenuation relationships for Sumatran-subduction earthquakes and the seismic hazard implications to Singapore and Kuala Lumpur, Soil Dynamics and Earthquake Engineering, 23, 779–788, https://doi.org/10.1016/S0267-7261(03)00072-3, 2003.

Mousavi, M., Pai, J. S., and Goda, K.: Designing parametric earthquake insurance: Trigger calibration and basis risk evaluation, Insurance: Mathematics and Economics, 110, 188–203, https://doi.org/10.1016/j.insmatheco.2023.02.005, 2023.

OECD: Enhancing Financial Protection Against Catastrophe Risks: The Role of Catastrophe Risk Insurance Programmes, Tech. rep., OECD Publishing, Paris, 2021.

Pagani, M., Garcia-Pelaez, J., Gee, R., Johnson, K., Poggi, V., Styron, R., Weatherill, G., Simionato, M., Viganò, D., Danciu, L., and Monelli, D.: The 2018 version of the Global Earthquake Model: Hazard component, Earthquake Spectra, 39, 2012–2037, https://doi.org/10.1177/87552930231180269, 2023.

Pai, J. S., Mousavi, M., and Tan, K. S.: Modelling parametric earthquake insurance with multi-trigger structures, Geneva Risk and Insurance Review, 47, 145–172, https://doi.org/10.1057/s10713-022-00067-w, 2022.

Reasenberg, P.: Second-order moment of central California seismicity, 1969–1982, Journal of Geophysical Research, 90, 5479–5495, https://doi.org/10.1029/JB090iB07p05479, 1985.

Salgado-Gálvez, M. A., Zuloaga, D., Velasquez, C. A., Carreño, M. L., Cardona, O.-D., and Bazzurro, P.: Development of a regionally consistent and fully probabilistic earthquake risk model for Central Asia, Natural Hazards and Earth System Sciences, 24, 3851–3868, https://doi.org/10.5194/nhess-24-3851-2024, 2024.

Scordilis, E. M.: Empirical global relations converting M_s and m_b to moment magnitude, Journal of Seismology, 10, 225–236, https://doi.org/10.1007/s10950-006-9012-4, 2006.

Silva, V., Crowley, H., Pagani, M., Monelli, D., and Pinho, R.: Development of the OpenQuake engine, the Global Earthquake Model's open-source software for seismic risk assessment, Natural Hazards, 103, 1969–1992, https://doi.org/10.1007/s11069-020-04127-0, 2020.

Storchak, D. A., Di Giacomo, D., Bondár, I., Engdahl, E. R., Harris, J., Lee, W. H. K., Villaseñor, A., and Bormann, P.: Public Release of the ISC-GEM Global Instrumental Earthquake Catalogue (1900–2009), Seismological Research Letters, 84, 810–815, https://doi.org/10.1785/0220130034, 2013.

Swiss Re Institute: Closing the protection gap: Disaster risk financing, Tech. rep., Swiss Re Sigma 5/2016, 2016.

Utsu, T.: Representation and analysis of the earthquake size distribution: A historical review and some new approaches, Pure and Applied Geophysics, 155, 509–535, https://doi.org/10.1007/s000240050276, 1999.

Wiemer, S. and Wyss, M.: Minimum magnitude of completeness in earthquake catalogs: Examples from Alaska, the western United States, and Japan, Bulletin of the Seismological Society of America, 90, 859–869, https://doi.org/10.1785/0119990114, 2000.

Woessner, J. and Wiemer, S.: Assessing the quality of earthquake catalogues: Estimating the magnitude of completeness and its uncertainty, Bulletin of the Seismological Society of America, 95, 684–698, https://doi.org/10.1785/0120040007, 2005.

World Bank: Indonesia: Disaster Risk Financing and Insurance — Country Note, Tech. rep., World Bank, Washington, DC, 2017.

Zaliapin, I. and Ben-Zion, Y.: Earthquake clusters in southern California I: Identification and stability, Journal of Geophysical Research: Solid Earth, 118, 2847–2864, https://doi.org/10.1002/jgrb.50179, 2013.

Zhao, J. X., Zhang, J., Asano, A., Ohno, Y., Oouchi, T., Takahashi, T., Ogawa, H., Irikura, K., Thio, H. K., Somerville, P. G., Fukushima, Y., and Fukushima, Y.: Attenuation relations of strong ground motion in Japan using site classification based on predominant period, Bulletin of the Seismological Society of America, 96, 898–913, https://doi.org/10.1785/0120050122, 2006.