

Monte Carlo Simulation of Earthquake Damage to Mexico City: A Cadastral Exposure Model with Voronoi-Weighted Spatio-Temporal Point Processes

Geronimo Rojas Barragan

P41722@correo.uia.mx

Universidad Iberoamericana: Universidad Iberoamericana <https://orcid.org/0009-0007-6415-4112>

Nelson Muriel Torrero

Universidad Iberoamericana

Research Article

Keywords: ETAS point process, Voronoi residuals, cadastral exposure, ground motion prediction equations, catastrophe modelling, Mexico City, Monte Carlo Simulation

Posted Date: September 30th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10944185/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

1 Monte Carlo Simulation of Earthquake Damage to
2 Mexico City:
3 A Cadastral Exposure Model with Voronoi-Weighted
4 Spatio-Temporal Point Processes

5 Gerónimo Rojas Barragán¹ and Nelson Omar Muriel Torrero²

6 ¹FISMAT, Universidad Iberoamericana

7 ²FISMAT, Universidad Iberoamericana

8 2026

9 **Abstract**

10 We estimate the monetary damage that future earthquakes would inflict on
11 Mexico City, using the cadastral register of the city — the government’s own
12 assessment of property value, covering both land and the constructions upon it —
13 as the exposure base. The model has three fitted layers. First, an epidemic-type
14 aftershock sequence (ETAS) point process is fitted separately to two seismic
15 sources recovered by k -means clustering of the Mexican catalogue on magnitude
16 and focal depth. Second, ground motion at the city is predicted by regime-specific
17 attenuation relations anchored to two observed firm-ground records, and amplified
18 by a site factor. Third, a discrete damage-state vulnerability model converts
19 ground motion into a monetary loss ratio against each cadastral unit. The
20 vulnerability layer is calibrated by joint maximum likelihood to *five* independent
21 observables of the 19 September 2017 M_w 7.1 Puebla–Morelos earthquake Monte
22 Carlo simulation of 2×10^5 independent centuries (2×10^7 simulated years)
23 per assumption bound gives an average annual loss of 2.26 bn MXN (bracket
24 0.14–27.0) and a 1-in-100-year annual loss of 40 bn MXN (bracket 1.4–439).

25 **Keywords:** ETAS point process; Voronoi residuals; cadastral exposure; ground motion
26 prediction equations; catastrophe modelling; Mexico City. Monte Carlo Simulation.

27 **Contents**

28	1 Introduction	2
29	1.1 What cadastral value measures, and what the 2021 extract contains . .	4
30	1.2 Structure of the paper	5

31	2 Data and catalogue completeness	5
32	2.1 The earthquake catalogue	5
33	2.2 The 1980 detection step	5
34	2.3 Cadastral exposure	6
35	3 Point-process model	8
36	3.1 Voronoi spatial weighting	8
37	3.2 Source clustering	9
38	3.3 Temporal model	11
39	4 From earthquake to damage	13
40	4.1 Ground motion	13
41	4.2 Vulnerability	13
42	4.3 Site amplification and the Atlas 2017 dataset	14
43	5 Calibration to the 2017 event	14
44	5.1 Why not a scaling constant	14
45	5.2 The statistical model	14
46	5.3 The reporting-completeness result	16
47	5.4 Goodness of fit	17
48	5.5 A warning on scope	21
49	6 Monte Carlo simulation	22
50	6.1 Engine	22
51	6.2 The simulation in one page	22
52	6.3 On the number of simulations	23
53	6.4 Results	23
54	6.5 Tail risk: value-at-risk and tail value-at-risk	23
55	7 Discussion and limitations	25
56	7.1 The bracket is the result	25
57	7.2 Expected expenditure over the next century, in 2026 pesos	26
58	7.3 Computational cost	28
59	7.4 Known limitations	28
60	7.5 Next steps	28
61	Reproducibility	29
62	Acknowledgements and tooling	29
63	8 References	29

1 Introduction

Mexico City exhibits exceptional seismic vulnerability due to site-response mechanisms that operate largely independently of local epicentral frequency. The metropolitan area

sits atop the saturated lacustrine deposits of a drained lake bed, located several hundred kilometers from the active subduction zone along the Pacific coast. These soft sediments strongly amplify long-period seismic waves, turning distant, moderate-to-large events into destructive catastrophes. This dynamic was dramatically demonstrated during both the 1985 Michoacán (M_w 8.0) earthquake—located 384 km from the capital—and the 2017 Puebla–Morelos (M_w 7.1) earthquake at a distance of 135 km. Both events induced severe structural collapse across the urban basin despite producing modest ground accelerations on firm ground outside the valley.

Building upon the spatial modeling framework of [14]—who utilized Voronoi deviance residuals and spatial Poisson processes to evaluate seismic hazards in Canada—we adapt and extend this methodology to the complex tectonic and socioeconomic context of Mexico. Our approach introduces two key extensions:

First, we account for distinct seismotectonic regimes across the Mexican landmass. While existing models often assume homogeneous point-process characteristics, our geographic partition captures multiple seismic sources, including the Neovolcanic Belt and subduction zones associated with the Cocos and Pacific plate interfaces (e.g., the San Andreas fault system).

Second, we incorporate a high-resolution exposure inventory by leveraging Mexico City’s municipal cadastral register. Unlike standard macroeconomic aggregates or regional population proxies, this dataset provides block-level cadastral values across all boroughs (alcaldías), representing one of the most granular exposure registers available for any global megacity.

Quantifying the monetary consequences of seismic risk requires integrating three fundamental components: a probabilistic seismicity model estimating event occurrence and magnitude distributions, a ground motion attenuation model capturing site amplification within the basin, and a precise inventory of assets at risk. The primary contribution of this paper is the unified construction of all three elements using empirical Mexican data, anchored by the cadastral register as the baseline financial exposure map.

The proposed framework proceeds in four sequential stages:

1. **Spatial Partitioning and Source Disaggregation:** We map historic catalog events across Mexico and construct a Voronoi tessellation to discretized spatial intensity, ensuring a single spatial centroid per polygon. To capture spatial heterogeneity across tectonic regimes, we apply k -means clustering to the tessellated events. This disaggregates historical seismicity into distinct source clusters, accounting for variations in depth, recurrence intervals, and magnitude distributions between intraslab, interplate, and crustal faulting events.
2. **Spatio-Temporal Point Process Simulation:** Following [14], we formulate a spatio-temporal Poisson process with a multiplicative intensity structure. This enables joint simulation of synthetic earthquake vectors (x, y, t, m) specifying spatial coordinates (x, y) , future occurrence time t , and moment magnitude m .
3. **Attenuated Financial Damage Function:** To map simulated hypocenters

to structural and monetary loss within the city, we implement an attenuation loss function parameterized by epicentral distance and basin amplification effects [3]. Given the unique amplification profile of the lacustrine basin, ground motion attenuation is non-linear and dictates disproportionate damage potential for specific structural resonant frequencies.

4. **Monte Carlo Aggregation:** Finally, we perform extensive Monte Carlo simulations across synthetic catalog histories to compute the empirical loss distribution and exceedance probability curves for Mexico City. High-performance computational execution was utilized to iterate over the spatial grid and process the large-scale cadastral matrix.

1.1 What cadastral value measures, and what the 2021 extract contains

Cadastral value is the reason this exposure base is worth building on, so it is worth being precise about what it covers and what the public release of it omits. We verified that the value column in the extract satisfies

$$\text{valor_suelo} = \text{valor_unitario_suelo} \times \text{sup_terreno}$$

exactly, with a ratio of 1.0000 across the distribution: median, 1st and 99th percentiles all 1.0000, and not one row of the 2,295,465 in which `valor_suelo` exceeds unit value times land area.

This requires care in interpretation. Cadastral value in Mexico City is *not* a land-only quantity. Article 129 of the *Código Fiscal de la Ciudad de México* requires the annual issuance of unit values for land, for the constructions attached to it, and for special installations and complementary works, and taxpayers combine these to determine the cadastral value of a property; the unit values are set to track market prices of land *and* construction. *Plusvalía* denotes the location premium — the same building being worth more in one place than another — and is embedded in the land unit values by design, not excluded from them. Cadastral value is thus the closest available proxy for gross property value, which is precisely what recommends it as an exposure base.

What is limited is the *public extract*, not the concept. The 2021 release used here publishes only the land component: its sole value fields are `valor_unitario_suelo` and `valor_suelo`, with no construction-value column, even though 2,268,879 records (98.8%) report `sup_construccion` > 0. The buildings are present in the data; only their assessed value is absent. Recovering the construction component requires the *Valores Unitarios de las Construcciones* tables of the Fiscal Code for the matching year, applied to `sup_construccion` by construction class. For fiscal year 2021 those tables were published in the *Gaceta Oficial* of 21 December 2020 as scanned images with no text layer, so they require manual transcription and are not machine-readable.

Pending that transcription, we use built floor area with an explicit replacement cost of 12,000 MXN/m², giving a structural exposure of 5.12 trillion MXN. Every monetary

146 result below is proportional to this constant, and we report it explicitly for that
 147 reason. This is a substitute for the cadastral construction value, not an improvement
 148 on it: the official figures are assessed per parcel and per construction class, whereas a
 149 single constant cannot distinguish a reinforced-concrete tower from a masonry dwelling.
 150 Replacing it is the highest-priority correction to the exposure layer (§7.4).

151 A second property required verification. The parcel identifier `fid` indexes a *predio* (a
 152 land parcel), and condominium units on the same parcel share it: 2,295,465 cadastral
 153 records join to 1,089,684 surveyed polygons, so the relation is many-to-one in fifteen of
 154 sixteen boroughs. Summing value over records could therefore double-count land. It
 155 does not. Testing against polygon areas in Benito Juárez,

$$\frac{\sum_i \text{sup_terreno}_i}{\text{polygon area}} = 1.009 \quad (\text{IQR } 0.984\text{--}1.101),$$

156 whereas taking one record per parcel gives 0.103. Each unit therefore carries its
 157 apportioned *indiviso* share, and record-level summation is correct.

158 1.2 Structure of the paper

159 Section describes the earthquake catalogue and establishes its magnitude of completeness.
 160 Section sets out the point-process model, the Voronoi spatial weighting, and the
 161 source clustering. Section sec:hazard gives the ground-motion and vulnerability layers.
 162 Section calibrates the vulnerability model to the 2017 event and presents the reporting-
 163 completeness result. Section presents the Monte Carlo engine and the century-scale
 164 loss projections. Section states the limitations, of which the missing microzonation is
 165 the most consequential.

166 2 Data and catalogue completeness

167 2.1 The earthquake catalogue

168 The catalogue is the Servicio Sismológico Nacional (SSN) record, 358,353 events from 20
 169 January 1900 to 13 October 2025. After removing 17,656 records whose magnitude field
 170 reads `no calculable`, seven with invalid coordinates, and fourteen duplicate epicentres,
 171 and applying a threshold of $M \geq 4.95$, 1,848 events remain (Figure 1).

172 2.2 The 1980 detection step

173 Fitting a point process to the full 1900–2025 span would be a serious error, and
 174 identifying why is a prerequisite for everything that follows. Table 1 shows the catalogue
 175 rate by era.

176 Two diagnostics settle the question. The rate of $M \geq 6.5$ events is stable across all
 177 four eras (1.35–2.00 yr^{−1}): large earthquakes were always detected, so the underlying
 178 seismicity is stationary. Meanwhile the *b*-value at $M_c = 5.0$ rises from 0.22 to 1.03. A




Figure 1: Seismicity of Mexico, $M \geq 4.95$, 1900–2025 (1,848 events after cleaning). Marker area scales with magnitude; colour encodes magnitude. The three largest events in the catalogue are labelled. Activity concentrates along the Pacific subduction trench and the Gulf of California transform system, with the northeast of the country effectively aseismic.

179 b -value near unity is the expected Gutenberg–Richter value; 0.22 is the signature of a
 180 catalogue missing its smaller events.

181 We therefore fit only the post-1980 record at $M_c = 5.0$. The consequence is not cosmetic.
 182 Fitting 1900–2025 at $M \geq 4.95$ yields a degenerate ETAS fit in which the productivity
 183 exponent collapses to $\alpha = 0.065$ with a Kolmogorov–Smirnov p -value of 0.35: the
 184 model absorbs the completeness step by attributing it to triggering. Restricted to 1980
 185 onward, the same model gives $\alpha = 1.037$ with $p = 0.999$ for the shallow source (Figure).
 186 The pre-1980 record remains usable at $M_c = 6.5$, where the whole 1900–2025 span is
 187 complete — which is the appropriate catalogue for estimating the recurrence of the
 188 largest events.

189 2.3 Cadastral exposure

190 Figure 3 maps unit land value across the city. The range is wide: median assessed
 191 value runs from 611 MXN/m² in Milpa Alta to 6,144 in Miguel Hidalgo, a factor
 192 of ten. Iztapalapa holds the most parcels (366324); Miguel Hidalgo the most value
 193 (186 billion MXN).

Era	n	Rate (yr^{-1})	min M	$b(M_c=5.0)$	$M \geq 6.5$ rate
1900–1970	116	1.66	6.30	0.22	1.61
1970–1980	186	18.60	5.00	0.86	2.00
1980–2000	620	31.00	5.00	1.02	1.90
2000–2026	926	35.62	5.00	1.03	1.35

Table 1: Catalogue completeness by era. The $M \geq 5$ rate rises eighteen-fold while the $M \geq 6.5$ rate is essentially constant, and the Gutenberg–Richter b -value at $M_c = 5.0$ rises from 0.22 to 1.03. Both signatures identify an instrumentation change, not a change in seismicity.




Figure 2: (a) The full magnitude–time record. Before 1980 (shaded) only $M \geq 6.3$ events appear; the $M = 6.5$ line is complete over the whole period. (b) ETAS conditional intensity for the shallow source on the complete window; spikes are aftershock sequences. (c) On the complete window the ETAS transformed time $\tau(t)$ tracks the observed count (KS $p = 0.999$) whereas a homogeneous Poisson process does not.




Figure 3: (a) Median cadastral unit land value on a 1 km hexagonal grid, quantile-classified, with the Atlas 2017 zona cero footprint dashed. (b) Median unit land value by alcaldía. Note that the highest-value districts (Miguel Hidalgo, Benito Juárez, Cuauhtémoc) lie substantially inside the observed 2017 damage footprint.

3 Point-process model

3.1 Voronoi spatial weighting

Following the Voronoi-residual approach, we tessellate the epicentre field and weight each cell by inverse area. For a tessellation $\{A_i\}$ over the study region,

$$A_T = \sum_i \frac{1}{A(i)}, \quad P_i = \frac{1}{A(i) A_T}, \quad \sum_i P_i = 1. \quad (1)$$

Dense clustering produces small cells and hence large weight, so P_i acts as a nonparametric estimate of the spatial intensity (Figure 4).

Two implementation points matter. Areas are computed on an equal-area projection rather than from latitude-longitude degrees: a degree of longitude spans 107 km at latitude 14 and 93 km at latitude 33, so unprojected areas are biased along precisely the axis on which the Mexican catalogue is elongated. Second, the study region is the Mexican landmass buffered 300 km offshore, not the landmass itself. Subduction epicentres lie offshore by construction; clipping to land would delete 1236 of 1848 cells, including most of the events that matter.

Dense clustering along the trench → small cells → high probability weight (1,770 cells, median 275 km²)




Figure 4: (a) Voronoi tessellation of the 1848 epicentres, clipped to the study region. (b) Inverse cell area, which is the spatial weight P_i up to normalisation. Median cell area is 275 km².

3.2 Source clustering

We partition the catalogue into two seismic sources by k -means on magnitude and focal depth, *without* standardising the features. This deserves explicit justification, since standardisation is the usual default.

Depth carries roughly 4,177 times the variance of magnitude in this catalogue, so an unstandardised metric is effectively depth-only. That is the physically correct behaviour here: the quantity that distinguishes Mexican seismic sources *is* depth — shallow interface and transform events versus deep events within the subducting Cocos plate. Standardising promotes magnitude to equal footing, and the algorithm then splits on magnitude instead, returning centres at M 5.24 and M 6.65 at nearly identical depths (35.6 and 41.2 km), which is a large/small split rather than a source partition.

The two options can be compared quantitatively. Measuring the distance between cluster centroids relative to within-cluster spatial spread:

Feature set	Centroid separation	Separation / spread
(magnitude, depth), raw	761 km	1.15
(magnitude, depth), standardised	92 km	0.10

The unstandardised clustering separates the populations roughly nine times better in space, and its silhouette score is higher (0.749 versus 0.564). Its deep cluster lies entirely within latitudes 13.0–19.9 (IQR 14.5–17.0); no Baja California event falls in it (mean depth there is 13 km), while 33% of Chiapas–Tehuantepec events do, where the slab is deepest. The partition is therefore geographically coherent (Figure 5).

226 This matters beyond taxonomy: the 40 km depth boundary between the two clusters is
 227 exactly the boundary at which the ground-motion model switches regime (Section 4.1).
 228 The clustering selects the attenuation relation.




Figure 5: The two seismic sources recovered by *k*-means on (magnitude, depth), plotted geographically. Deep intraslab events (orange) are confined to the southern trench and the Isthmus; the Baja California transform system contains none, so the shallow interface cluster alone supplies the northern seismicity. Marker area scales with magnitude. Catalogue 1980–2025, $M \geq 5.0$, $n = 1,848$.




Figure 6: The same partition in feature space. The clusters separate cleanly on depth: the deepest shallow-cluster event is at 65 km and the shallowest deep-cluster event at 66 km, so the k -means boundary is effectively a cut at 65.5 km. The 40 km line is the separate, physically motivated threshold that selects the ground-motion prediction equation for each individual event; the cluster label sets which ETAS process generates the event, not which attenuation relation applies to it.

3.3 Temporal model

Each source is modelled as an ETAS process with conditional intensity

$$\lambda(t \mid \mathcal{H}_t) = \mu + A \sum_{t_i < t} e^{\alpha m_i} \left(1 + \frac{t - t_i}{c}\right)^{-p}, \quad m_i = M_i - M_0 \geq 0, \quad (2)$$

fitted by maximum likelihood over $[0, T]$ with the closed-form integrated intensity. Table 2 gives the fitted parameters.

Source	n	μ (d ⁻¹)	A	α	c	p	KS p
Shallow interface	1288	0.0362	1.5379	1.037	0.0127	0.978	0.999
Deep intraslab	258	0.0132	0.2035	1.535	0.0114	0.898	0.413

Table 2: ETAS parameters by source, fitted to the completeness-corrected catalogue (1980–2025, $M_c = 5.0$), exponential magnitude marks.

233 The Omori kernel and its integral

234 The temporal factor $(1 + (t - t_i)/c)^{-p}$ in the intensity above is the *modified Omori*
 235 *formula*, in the normalised parameterisation (the kernel equals unity at $t = t_i$, so the
 236 productivity A carries the units). Omori [15] observed that the rate of aftershocks
 237 following a mainshock decays as the reciprocal of elapsed time, $n(t) \propto (t + c)^{-1}$; the
 238 generalisation to a free exponent, $n(t) = K(t + c)^{-p}$ — the *modified Omori formula*,
 239 conventionally attributed to Utsu [16] — is the form in use today. Utsu et al. [17]
 240 review a century of fits to observed sequences and report that p typically falls in the
 241 range 0.9–1.5, varying from sequence to sequence with no systematic dependence on
 242 mainshock magnitude. In the ETAS formulation of Ogata [18] this decay law is not
 243 applied to a single designated mainshock: every event, whatever its size, contributes
 244 its own Omori kernel to the conditional intensity, weighted by $e^{\alpha m_i}$. Aftershocks are
 245 thereby produced by the same mechanism as mainshocks, and no event needs to be
 246 labelled in advance — which is precisely the property this study requires, since the
 247 catalogue arrives unlabelled and the simulation must generate cascades without being
 248 told which events are triggers.

249 The quantity that matters computationally is the *integral* of that kernel, because it
 250 appears twice: once in the log-likelihood, where the compensator $\int_0^T \lambda(t) dt$ must be
 251 evaluated at every optimiser step, and once in the simulation, where the expected
 252 number of direct offspring of an event is $Ae^{\alpha m_i}$ times the kernel integral over the
 253 remaining horizon. It has a closed form,

$$\int_0^T \left(1 + \frac{u}{c}\right)^{-p} du = \begin{cases} \frac{c}{p-1} \left[1 - \left(1 + \frac{T}{c}\right)^{1-p}\right], & p \neq 1, \\ c \log\left(1 + \frac{T}{c}\right), & p = 1, \end{cases} \quad (3)$$

254 which is why the likelihood needs no numerical quadrature and the simulation can
 255 draw offspring counts directly from a Poisson distribution rather than by thinning.
 256 Inverting the same expression gives the offspring delays by exact inverse-CDF sampling,
 257 so no rejection step is needed either — this is the single reason a century of cascade
 258 generation costs a fraction of a millisecond (§7.3).

259 Equation (3) also makes the consequence of $p < 1$ explicit. Our fits give $p = 0.978$ and
 260 $p = 0.898$ — the first inside the range reported by [17], the second marginally below it
 261 — and for $p < 1$ the exponent $1 - p$ is positive, so $(1 + T/c)^{1-p}$ grows without bound as
 262 $T \rightarrow \infty$. The integral therefore diverges and the nominal branching ratio is unbounded.
 263 This is a statement about an infinite-time idealisation, not about the physical process:
 264 over any bounded horizon the integral is finite. Evaluating it over 100 years gives
 265 effective branching ratios of 0.569 and 0.156, both comfortably subcritical, so cascades
 266 terminate and the expected event count per realisation is finite. Bounded-horizon
 267 simulation is therefore well posed, and we never simulate open-endedly. The alternative
 268 — constraining $p > 1$ during fitting to force convergence at infinity — would impose a
 269 bound the data do not support in order to make an idealisation behave, and would
 270 degrade the fit at the timescales that actually carry loss.

271 4 From earthquake to damage

272 The link from an earthquake at position (x, y) with magnitude M and depth h to a
 273 monetary loss in Mexico City is a chain of factors, each of which carries an explicit
 274 bracket:

$$L(x, y, M, h) = \underbrace{S_{\text{rock}}(M, R)}_{\text{attenuation}} \times \underbrace{F_{\text{site}}}_{\text{amplification}} \times \underbrace{e^{\varepsilon}}_{\text{aleatory}} \longrightarrow \underbrace{\text{MDR}}_{\text{fragility}} \times \underbrace{V}_{\text{cadastral value}}. \quad (4)$$

275 4.1 Ground motion

276 We use the functional form

$$\ln \text{PGA} = c_0 + c_1(M - 6) + c_2 \ln R_{\text{hyp}} + c_3 R_{\text{hyp}}, \quad (5)$$

277 with separate coefficients for the interface and intraslab regimes, split at 40 km depth.
 278 Magnitude and distance slopes are set to values typical of Mexican attenuation, which
 279 is stronger than global relations [1]; intercepts are solved so that each regime reproduces
 280 an observed firm-ground record at Ciudad Universitaria (seismic Zone I):

	Event	Regime	R_{hyp}	Observed	Model
281	1985-09-19 M_w 8.0	interface	384 km	34 cm/s ²	38.0
	2017-09-19 M_w 7.1	intraslab	135 km	57 cm/s ²	55.5
	2017-09-07 M_w 8.2	intraslab	739 km	15 cm/s ²	5.5

282 The third row is an out-of-sample test and the model under-predicts it by a factor of
 283 three. The functional form loses accuracy beyond roughly 500 km, and we report this
 284 rather than restricting the test set. For publication the anchored intercepts should be
 285 replaced by the published coefficient tables of [1] for interface events and [2] — or its
 286 2017-inclusive update [3] — for intraslab events.

287 4.2 Vulnerability

288 A single lognormal mean-damage-ratio curve cannot reproduce the 2017 event. Any
 289 threshold matching the aggregate loss spreads light damage across roughly 14% of the
 290 building stock when only 0.52% was damaged; forcing both drives the threshold to
 291 physically meaningless values (3097–7681 cm/s², i.e. 50% damage at three to eight g).
 292 Portfolio loss at moderate shaking is driven by a small fraction of badly damaged
 293 buildings, not by every building being slightly damaged.

294 We therefore use discrete damage states with consequence ratios. For state ds with
 295 threshold θ_{ds} and dispersion β ,

$$P(ds) = \Phi\left(\frac{\ln(\text{PGA}/\theta_{ds})}{\beta}\right), \quad \text{MDR} = \sum_{ds} [P(ds) - P(ds+1)] r_{ds}, \quad (6)$$

296 with consequence ratios $r = (0.02, 0.10, 0.50, 1.00)$ for slight, moderate, extensive and
 297 complete damage. Thresholds scale by construction era, with multipliers (1.00, 1.55, 2.30)
 298 for pre-1985, 1985–2004 and post-2004 construction — the two revisions of the city’s
 299 building code that these earthquakes precipitated.

300 4.3 Site amplification and the Atlas 2017 dataset

301 Site amplification is the largest single factor in Equation 4 and the least constrained by
302 the available data. The appropriate input is the NTC-Sismo microzonation (Zones I, II,
303 IIIa–IIIId), which is not in the project’s data holdings.

304 The Atlas de Riesgo 2017 *zona cero* dataset [12] is sometimes treated as a substitute.
305 It is not a microzonation: it comprises 22 polygons totalling 348.6 km², attributed only
306 by administrative reconstruction quadrant, and carries no damage values. It covers
307 59.6% of cadastral parcels and 49.1% of assessed land value.

308 What it does encode is where damage actually concentrated in an observed M_w 7.1
309 event, which is empirically valuable even though it is not a soil model. We use it in
310 two ways: as the observed damaged footprint, one of the four calibration targets; and
311 as a proxy for the lacustrine/transition site split, assigning amplification factors of 4.8
312 inside and 1.5 outside. Substituting the real microzonation is the single highest-value
313 improvement available to this model, and Section 7 quantifies why.

314 5 Calibration to the 2017 event

315 5.1 Why not a scaling constant

316 The simplest calibration would divide observed loss by modelled loss and multiply
317 everything by the ratio. We reject this for three reasons.

318 First, the 2017 event is observed through five independent quantities, and a single
319 multiplicative constant forces one of them to match while leaving the other four free.
320 Fitting one observable with one parameter is not evidence that a model reproduces an
321 event.

322 Second, a constant is a point estimate with no standard error, so the reported loss
323 bracket would have to be assumed rather than derived.

324 Third, and most importantly, a constant cannot fail. A joint fit of a small number of
325 parameters to five observables retains degrees of freedom: if the model form is wrong, it
326 *cannot* satisfy all five simultaneously, and the residuals reveal it. That is a genuine test.

327 5.2 The statistical model

328 Let $\psi \in \Theta \subseteq \mathbb{R}^4$ denote the vector of structural vulnerability parameters, defined as:

$$\psi \doteq \begin{pmatrix} \theta_c \\ \beta \\ s_c \\ \rho \end{pmatrix} \in (0, \infty) \times (0, \infty) \times (0, 1] \times (0, 1], \quad (7)$$

329 where:

- 330 • $\theta_c \in (0, \infty)$ represents the complete-damage acceleration threshold,

- $\beta \in (0, \infty)$ denotes the lognormal dispersion (fragility curve slope parameter),
- $s_c \in (0, 1]$ parameterizes the spatial discretization interval (spacing) of the complete damage state, and
- $\rho \in (0, 1]$ is the empirical reporting completeness ratio.

For a given parameter realization $\boldsymbol{\psi}$, let $\mathcal{M}_{\boldsymbol{\psi}} : \Theta \rightarrow \mathbb{R}^K$ be the deterministic forward damage model. The model yields a set of K predicted observable values $\hat{\mathbf{y}}(\boldsymbol{\psi}) \doteq (\hat{y}_1(\boldsymbol{\psi}), \hat{y}_2(\boldsymbol{\psi}), \dots, \hat{y}_K(\boldsymbol{\psi}))^T$, where each component $\hat{y}_k(\boldsymbol{\psi})$ corresponds to the k -th physical observable recorded during the September 19, 2017 $M_w 7.1$ Puebla–Morelos earthquake sequence, for $k \in \{1, 2, \dots, K\}$.

Let $\mathbf{y}^{\text{obs}} = (y_1^{\text{obs}}, \dots, y_K^{\text{obs}})^T$ denote the vector of realized empirical observations. Under the assumption of conditional independence given $\boldsymbol{\psi}$, the full log-likelihood function $\ell(\boldsymbol{\psi}; \mathbf{y}^{\text{obs}})$ is constructed as the additive sum of individual marginal log-likelihood contributions:

$$\ell(\boldsymbol{\psi}; \mathbf{y}^{\text{obs}}) = \sum_{k=1}^K \ln f_k(y_k^{\text{obs}}; \hat{y}_k(\boldsymbol{\psi}), \boldsymbol{\phi}_k), \quad (8)$$

where $f_k(\cdot; \hat{y}_k(\boldsymbol{\psi}), \boldsymbol{\phi}_k)$ represents the probability density (or probability mass) function governing the k -th observational error process, uniquely parameterized by the deterministic model forecast $\hat{y}_k(\boldsymbol{\psi})$ and a vector of observable-specific nuisance error parameters $\boldsymbol{\phi}_k$.

$$\ell(\boldsymbol{\psi}) = \sum_{k=1}^K \ell_k(y_k | \hat{y}_k(\boldsymbol{\psi})), \quad (9)$$

with $K = 5$ and the five terms as follows. Monetary loss and mapped footprint are positive, multiplicatively uncertain quantities, so they take a lognormal term

$$\ell_k = -\log(y_k \sigma_k \sqrt{2\pi}) - \frac{(\log y_k - \log \hat{y}_k)^2}{2\sigma_k^2}, \quad (10)$$

with σ_k the assumed log-scale spread of the published figure. The collapse count is a census of rare, effectively independent structural failures, so it takes a Poisson term

$$\ell_k = y_k \log \hat{y}_k - \hat{y}_k - \log \Gamma(y_k + 1). \quad (11)$$

The severely-damaged count is an *inspection* outcome rather than a census: inspection effort clusters spatially, so the count is overdispersed relative to Poisson and takes a negative-binomial term with dispersion r ,

$$\ell_k = \log \Gamma(y_k + r) - \log \Gamma(r) - \log \Gamma(y_k + 1) + r \log \frac{r}{r + \hat{y}_k} + y_k \log \frac{\hat{y}_k}{r + \hat{y}_k}. \quad (12)$$

Finally, the pre-1985 share of collapsed buildings is a proportion out of the $n = 44$ collapses, so it takes a binomial term

$$\ell_k = y_k \log \hat{y}_k + (n - y_k) \log(1 - \hat{y}_k), \quad (13)$$

with $\hat{y}_k$ the modelled pre-1985 share. The estimator is $\hat{\psi} = \arg \max_{\psi} \ell(\psi)$, obtained by Nelder–Mead on $\log \psi$ from a grid of starting values so that positivity is enforced without constraints. Interval estimates come from the profile likelihood: for a scalar component ψ_j ,

$$\ell_{\text{prof}}(\psi_j) = \max_{\psi_{-j}} \ell(\psi_j, \psi_{-j}), \quad \text{CI}_{1-\alpha} = \left\{ \psi_j : 2[\ell(\hat{\psi}) - \ell_{\text{prof}}(\psi_j)] \leq \chi^2_{1-\alpha, 1} \right\}. \quad (14)$$

We use the profile rather than a Hessian approximation deliberately. The observed information at $\hat{\psi}$ has eigenvalues spanning 3 orders of magnitude and pairwise log-scale correlations up to 0.70, so the quadratic approximation is poor in the directions that matter; the log-scale standard error for ρ is 0.4, which a symmetric interval cannot represent on a parameter confined to $(0, 1]$. Figure 8 shows the four profile curves.

5.3 The reporting-completeness result

Fitting without ρ — asserting that all physical damage enters the official tally — reproduces the collapse count (43.7 against 44) and the pre-1985 share (0.92 against 0.91) but over-predicts monetary loss by a factor of 2.36 (37.7 bn MXN against 16.0 bn MXN) and under-predicts the footprint by 27%. It is rejected against the adopted model (deviance 6.92 on 1 df, $p = 0.009$; Table 4). Decomposing the modelled loss by damage state explains the direction: most of it is slight-to-moderate damage spread thinly across a 5.12 trillion MXN building stock, and that is precisely the component an official tally does not capture.

The published 16 bn MXN figure is a tally of *reported and inspected* damage — buildings surveyed, tagged, claimed on insurance, or entered into a reconstruction programme [11]. Cosmetic cracking that an owner repaired without reporting is real economic loss but appears in no official total. We therefore introduce ρ , the fraction of slight damage reaching the reported total, with moderate damage scaled proportionately but capped at unity (encoding that more severe damage is more likely to be counted). This is the same structure as a detection probability in capture–recapture or an under-reporting parameter in an epidemic model. At the fitted value, 64% of total physical loss is unreported.

The structure of the fit gives ρ a distinct role rather than making it a free scaling on the total: the collapse count constrains the upper tail, the footprint constrains spatial extent, and only the monetary observable constrains the reported fraction of the lower states. A pure scaling constant could not satisfy the count and area observables simultaneously, and setting ρ cannot be removed without cost: the era ratio and the reporting fraction act on different observables, and Table 4 shows that fixing the era ratio at its code-implied value instead is rejected ($p = 0.007$).

That is a statement about the *role* of ρ , not about how precisely one event determines it, and the two should not be conflated. With five observables ρ is bounded on both sides — its 95% profile interval is 0.052–0.599 (Figure 8c) — which is a material improvement on the four-observable fit, where the interval extended to the lower parameter bound

and severe under-reporting could not be excluded at all. It remains the least precisely determined of the four parameters, spanning an order of magnitude, and any conclusion depending on its exact value — including the split between reported and total economic loss — inherits that width. An independent estimate of reporting completeness, from insurance-claim or reconstruction-programme records, would let ρ be fixed externally rather than fitted, and is the single most valuable addition to this calibration.

5.4 Goodness of fit

The calibration uses five observables against four fitted parameters, leaving one degree of freedom and therefore a goodness-of-fit test the model can fail. It does not fail it: the deviance against the saturated model is 0.44 on 1 df ($p = 0.51$), and the four-parameter model has the lowest AIC of the three considered (Table 4). The residuals in Table 3 are consequently informative rather than forced — three observables are matched to within 3%, while the severe-damage count is over-predicted by $1.41\times$, which is the model’s one visible tension and is discussed below.

Two design decisions make the test possible, and both are worth stating because the obvious alternatives destroy it.

First, **the era effect is fitted, not assumed**. Earlier drafts of this model fixed the ratio of the complete-damage threshold between construction eras at a value read off the code revisions. That is a strong assumption about precisely the quantity a seismic-risk model most needs to get right, and it is also unnecessary, because [9] report the fraction of collapsed buildings predating 1985. Treating the era ratio k as a free parameter and that fraction as an observable turns an assumption into an estimate: $k = 1.297$ with a 95% profile interval of 1.200–1.451. Post-1985 construction is thus about 30% more resistant, in threshold terms, than pre-1985 construction — appreciably *less* of an improvement than the value 1.55 the code history alone would suggest, and the difference is statistically resolved rather than asserted.

Second, **the conventional parameter is fixed and the physical one is freed**. The damage-state spacing s_c is a bookkeeping convention of the damage-state formalism; the era ratio is a physical property of the building stock. Fixing $s_c = 1.0$ at its conventional value and freeing k yields a model with one degree of freedom that passes its test. Doing the reverse — as an earlier version of this analysis did — yields a saturated model that reproduces every observable exactly and tests nothing. Which parameter one chooses to fix is therefore not a matter of taste: fixing the convention buys a falsifiable model, and fixing the physics buys an unfalsifiable one.

The one tension is the severe-damage count, over-predicted by $1.41\times$ (848 modelled against about 600 reported). The direction is interpretable. “Severely damaged” is an inspection outcome, and inspection in the weeks after 19 September 2017 concentrated on buildings that were visibly compromised or whose occupants pressed for assessment; buildings damaged but occupied were systematically less likely to be tagged. A model that predicts physical damage should therefore over-predict a tally built from inspections, which is the sign observed. We do not add a second reporting parameter to absorb

436 it: that would restore saturation and discard the test, in exchange for a residual that
 437 already has a physical reading.

Observable	Observed	Model	Ratio	Error model
Reported loss	16.0 bn MXN	15.9 bn MXN	0.995	lognormal
Buildings collapsed	44	43.4	0.985	Poisson
Buildings severely dam.	600	848	1.41	negative binomial
Damaged footprint	348.6 km ²	339 km ²	0.972	lognormal
Pre-1985 share of collapses	0.91	0.91	0.999	binomial ($n = 44$)

Table 3: Calibration to the 19 September 2017 M_w 7.1 event. Five observables against four fitted parameters leaves one degree of freedom, so the residuals are informative: the deviance against the saturated model is 0.44 on 1 df ($p = 0.51$). Fitted parameters with 95% profile-likelihood intervals: $\theta_c = 1309$ cm/s² (1140–1477), $\beta = 0.419$ (0.378–0.453), $k = 1.297$ (1.200–1.451), $\rho = 0.117$ (0.052–0.599). The complete-state spacing is fixed at its conventional value $s_c = 1.0$. Log-likelihood -10.85 .

438 The fitted $\rho = 0.117$ implies that only 11.7% of slight damage and 63.6% of moderate
 439 damage entered the official tally, and hence that the total economic loss of the 2017
 440 event exceeded the reported 16 bn MXN. The point estimate is 44.4 bn MXN, a reported
 441 fraction of 0.359 (Figure 7d), with a 95% joint-likelihood interval of 37.4–51.4 bn.

442 Two cautions attach to that number. First, the interval propagates parameter uncer-
 443 tainty but not model uncertainty: it assumes the damage-state and consequence-ratio
 444 structure is correct. Second, ρ governs how loss divides between reported and unreported
 445 rather than how much physical damage occurs, so the total is comparatively robust to
 446 it while the *split* is the poorly determined quantity. The claim that total economic loss
 447 substantially exceeded the reported figure is well supported; the specific factor of 2.8
 448 carries the uncertainty of ρ , whose 95% interval spans 0.052–0.599.

449 A further result deserves emphasis. The model reproduces the 349 km² damaged
 450 footprint *from the attenuation relation alone*: no separate area model is fitted. This
 451 answers a question the model was not explicitly built to answer, namely how damage
 452 area scales with magnitude. A larger event does not merely deepen damage where 2017
 453 caused it; it expands the footprint, because the ground-motion contour corresponding
 454 to a given damage ratio moves outward with magnitude. Both effects follow from the
 455 distance dependence of the attenuation relation.

Model	$\dim(\psi)$	df	ℓ	AIC	p (LRT)
s_c fixed, k free (adopted)	4	1	-10.85	29.70	0.51
k fixed at 1.55, s_c free	4	1	-14.28	36.55	0.007
$\rho = 1$ (all damage reported)	3	2	-14.31	34.63	0.009
Both s_c and k free (saturated)	5	0	-10.63	31.26	—

Table 4: Model comparison. The likelihood-ratio test is against the saturated model on 1 df. The adopted model — conventional damage-state spacing fixed, era ratio fitted — is not rejected ($p = 0.51$) and has the lowest AIC. Fixing the era ratio at the value implied by the code history and freeing the spacing instead is rejected ($p = 0.007$): the assumed era ratio of 1.55 is inconsistent with the observed pre-1985 share of collapses. The saturated model reproduces every observable exactly and therefore supplies no test, which is why it is reported for comparison only and not adopted.

Observable	Observed	MLE	95% prediction interval	Error model
Loss (bn MXN)	16.0	15.9	7.0 – 29.2	lognormal
Buildings collapsed	44	43.4	30 – 59	Poisson
Buildings severely damaged	600	848	637 – 1072	negative binomial
Footprint (km ²)	348.6	339	251 – 477	lognormal
Pre-1985 share of collapses	0.91	0.91	0.81 – 0.96	binomial ($n = 44$)

Table 5: Parameter uncertainty propagated to each observable. Intervals are obtained by sampling the joint 95% likelihood region ($2\Delta\ell \leq \chi^2_{0.95,4} = 9.49$) and evaluating the model at each draw. Four of the five intervals cover the observation; the severe-damage count does not, which is the model’s one genuine tension (§5.4). Monetary loss is pinned to a factor of about 4.2, against a factor of 3.5 before the two additional observables were introduced.




Figure 7: (a) Modelled 2017 damage ratio on the cadastral grid, with the observed Atlas 2017 footprint dashed; the model recovers the footprint from the attenuation relation alone. (b) Model/observed ratio for all five observables; the severely-damaged count (orange) is the one the model over-predicts, the single genuine tension discussed in the text. (c) Profile likelihood for the complete-damage threshold, with the 95% interval shaded. (d) Loss by damage state, total economic versus reported: light damage is substantially under-reported.




Figure 8: (a–d) Profile likelihood for each vulnerability parameter, plotted as deviance $2[\ell(\hat{\psi}) - \ell_{\text{prof}}]$ against the parameter value; the dashed line is $\chi^2_{0.95,1} = 3.84$ and shading marks the resulting 95% interval. All four parameters are bounded on both sides, which the four-observable calibration did not achieve for the reporting fraction ρ (c): adding the severe-damage count and the pre-1985 share of collapses is what closes its lower bound. The era ratio k (d) is estimated at 1.297 (1.200–1.451), excluding the value 1.55 implied by the code history alone. (e) Model prediction divided by observation for each of the five observables, with joint 95% prediction intervals; four intervals cover the observation and the severe-damage count does not (§5.4).

5.5 A warning on scope

Figures for “the 2017 earthquake” differ by a factor of about 3.5 depending on scope. The 16 bn MXN figure is CDMX only, direct structural damage. Widely quoted USD 2–4 billion figures are national totals for the whole September 2017 sequence, spanning several states and mixing direct damage with indirect losses. Applying the national figure to a CDMX-only exposure base overstates vulnerability by that factor. We use the narrower and correct quantity, and recommend that any comparison state its scope explicitly.

464 6 Monte Carlo simulation

465 6.1 Engine

466 6.2 The simulation in one page

467 Stripped to its essentials, the engine is the following loop. It is stated here as a “proto-
468 algorithm” because every line corresponds to one concrete step in the code, and the
469 whole century-scale projection is nothing more than this loop run 2×10^5 times per
470 assumption bound.

for each of N simulated centuries:

for each source cluster s (shallow interface, deep intraslab):

- (1) **Choose a location.** Draw an epicentre from cluster s with probability $P_i \propto 1/A_i$, the inverse Voronoi-cell area of the historical epicentres — dense trench zones are more likely.
- (2) **Forecast the earthquakes.** Generate occurrence times and magnitudes from the fitted point process: background events as a *Poisson* process at rate μ_s , offspring from the Omori kernel; magnitudes from the tapered Gutenberg–Richter law.
- (3) **Attenuate.** Map each event to firm-ground motion at the city by the multiplicative relation $\text{PGA} = e^{c_0 + c_1(M-6) + c_2 \ln R + c_3 R}$, then multiply by the site-amplification factor F_{site} .
- (4) **Value the loss.** Convert soil motion to a mean damage ratio through the *log-normal* fragility curve, multiply by the cadastral value of each exposure cell, and sum over the city.
- (5) **Save** the per-year loss for this century.

end

end; then aggregate the N century histories into the annual-loss distribution, the exceedance curve, and the risk measures below.

471 Steps (1)–(2) are the point-process layer of Section 3; step (3) is the ground-motion
472 model of Section 4.1; step (4) is the vulnerability model of Section 5. The three
473 reductions described next make this loop fast enough to run at 2×10^7 simulated years
474 per bound without changing its result.

475 Each simulated future is built in three layers. Occurrence times and magnitudes come
476 from the fitted ETAS process for each source, generated by branching (background
477 events as a Poisson process, offspring cascades from the Omori kernel) rather than
478 thinning — exact, and fully vectorised. Epicentres are resampled from the observed
479 epicentres of the same cluster, preserving the real geometry of the interface and the
480 slab. Losses are then evaluated against the cadastral grid.

481 Two reductions make the computation tractable, both exact or measured rather than
482 assumed. Exposure is aggregated onto a 1 km grid crossed with construction era and
483 site zone, reducing 2,295,465 units to 3472 cells with total value preserved exactly; since
484 every source is more than 100 km away, within-cell distance variation is far below the

attenuation model’s own dispersion. Second, and more consequentially, note that the aleatory ground-motion term is *exactly* equivalent to a magnitude shift:

$$e^{c_0+c_1(M-6)+c_2 \ln R+c_3 R} \cdot e^\varepsilon = e^{c_0+c_1((M+\varepsilon/c_1)-6)+c_2 \ln R+c_3 R}. \quad (15)$$

The loss layer therefore collapses to a two-dimensional lookup over (epicentre, effective magnitude), verified against exact evaluation to a relative difference of 3×10^{-15} . Interpolation error on material losses is below 0.12% across seven orders of magnitude. The result is 0.40 ms per simulated century, against 5.0 s for direct evaluation.

Finally, events below $M = 5.5$ are screened out. This is measured, not assumed: in a representative century, events above $M = 5.5$ are 25% of the catalogue but carry 99.88% of the loss.

6.3 On the number of simulations

We run 2×10^5 independent centuries per assumption bound, i.e. 2×10^7 simulated years each. It is worth being explicit about why this is enough, since far larger runs are computationally within reach.

The standard error of a return-period estimate scales as $1/\sqrt{N_{\text{years}} \times P}$, so 2×10^7 simulated years determines the 1-in-1000-year loss to within a few percent. That is already an order of magnitude tighter than the spread induced by parameter uncertainty, which spans a factor of about 12 between the low and high bounds. Additional simulations buy precision in the tail; they cannot buy accuracy in a model conditioned on one event, however many observables that event supplies. Effort is better spent narrowing the vulnerability parameters — above all by obtaining the microzonation.

6.4 Results

Bound	AAL	RP100	RP250	RP500	Century total	$P(\text{zero yr})$
Low	0.140	1.36	5.13	12.05	14.0	0.157
Central	2.261	40.25	100.49	180.71	226.1	0.052
High	27.029	438.50	769.09	1085.28	2702.9	0.012

Table 6: Monte Carlo results, billions of MXN. AAL is average annual loss; RPx is the annual loss exceeded with probability $1/x$; century total is the mean cumulative 100-year loss. Each row is 2×10^7 simulated years. Bounds are the 95% profile-likelihood interval of the vulnerability parameters, not assumed percentages.

6.5 Tail risk: value-at-risk and tail value-at-risk

For an actuarial reader the annual-loss distribution is summarised most directly by its tail measures. Let L be the annual cadastral loss. The *value-at-risk* at level α is the quantile $\text{VaR}_\alpha = \inf\{\ell : \Pr(L \leq \ell) \geq \alpha\}$, and the *tail value-at-risk* (expected shortfall)

510 is the average loss in that tail, $\text{TVaR}_\alpha = \mathbb{E}[L \mid L \geq \text{VaR}_\alpha]$. By construction $\text{VaR}_{0.99}$
511 equals the 1-in-100-year return loss reported above; TVaR is the coherent measure that
512 additionally captures how heavy the tail is beyond that point, and it is TVaR, not VaR,
513 that a solvency capital calculation requires. Table 7 gives both at the 99, 99.5 and 99.9
514 percent levels, estimated directly from the 2×10^7 simulated years per bound.

Table 7: Value-at-risk and tail value-at-risk of the annual cadastral loss, billions of MXN, by assumption bound. $\text{VaR}_{0.99}$ coincides with the RP100 loss; TVaR is the mean loss beyond the corresponding VaR.

Bound	$\alpha = 0.99$		$\alpha = 0.995$		$\alpha = 0.999$	
	VaR	TVaR	VaR	TVaR	VaR	TVaR
Low	1.36	12.47	3.78	22.67	25.71	75.99
Central	40.25	142.99	81.43	229.40	301.40	559.72
High	438.50	846.25	677.78	1153.33	1452.51	1990.08

515 At the central bound the 99.5% tail value-at-risk is 229 bn MXN, roughly 14 times the
516 observed 2017 direct loss, and it grows to 560 bn MXN at the 99.9% level. The gap
517 between VaR and TVaR at each level — for instance 81 bn MXN versus 229 bn MXN
518 at $\alpha = 0.995$ — measures the weight of the extreme tail that a return-period statistic
519 alone conceals.

520 The central estimate of average annual loss is $2.26 \text{ bn MXN} \pm 0.005$ (Monte Carlo
521 standard error), or roughly 0.044% of the structural exposure per year. The 1-in-100-
522 year annual loss is 40 bn MXN, about two and a half times the reported 2017 loss —
523 which is consistent, since 2017 was a significant but not extreme event for the city. In the
524 central case 5.2% of years incur no loss at all; the median year sees only 0.013 bn MXN,
525 because the loss distribution is dominated by rare large events rather than by a steady
526 background.

527 The mean worst-single-event magnitude per century is M 7.13, close to the 2017 event,
528 and the distribution of century-worst losses (Figure 9e) shows that a repeat of 2017 sits
529 near the middle of what a century typically produces — not at its tail.




Figure 9: (a) Annual cadastral damage: six sample realisations (grey) with across-realisation quantiles; the 2017 reference is marked. (b) Loss exceedance curve for the three assumption bounds, truncated at the return period supported by the simulation size. (c) Return-period losses. (d) Distribution of 100-year cumulative loss. (e) Worst single event per century, by magnitude and loss.

7 Discussion and limitations

7.1 The bracket is the result

The low-to-high span of a factor of about 77 is not a defect of presentation; it is the honest width of what the available data support, and it should be reported as the result rather than the central value alone. Its dominant contributor is identifiable. Site amplification enters Equation 4 multiplicatively and is currently proxied from an administrative polygon set rather than measured, contributing a factor of roughly 3.2 on its own. Obtaining the NTC-Sismo microzonation would narrow the interval more than any other single change, including more simulations, better attenuation coefficients, or

finer exposure resolution.

7.2 Expected expenditure over the next century, in 2026 pesos

Table 8 gives the simulated expected loss by decade, in constant 2026 MXN. These are *real* figures: no inflation adjustment is applied, so they answer the question “what would this damage cost at today’s prices” rather than “what nominal sum will change hands”. A user who wants nominal pesos must apply an inflation path of their own choosing on top.

Two distinct intervals appear in the table and they must not be confused. The LCI/UCI columns are the 95% confidence interval on the *expected value*, from the Monte Carlo standard error $1.96 \hat{\sigma}/\sqrt{N}$ with $N = 2 \times 10^5$ realisations per bound; they are narrow (about $\pm 1.9\%$ of the mean at decade resolution) because that is simply how precisely 2×10^7 simulated years pin down a mean. They say nothing about how much any particular decade will cost. The 95th-percentile column is the *predictive* quantity: one decade in twenty exceeds it. It is roughly 4.0 times the mean, because the loss distribution is dominated by rare large events — the mean is not a typical outcome, and in 21–32% of simulated years the loss is exactly zero.

Decade	Low		Central					High	
	mean	med.	LCI	mean	UCI	med.	p95	mean	med.
2026–2035	16.25	0.039	54.43	55.56	56.70	2.09	214.06	194.0	34.84
2036–2045	17.44	0.058	59.44	60.61	61.78	2.73	244.89	213.0	41.56
2046–2055	18.47	0.067	62.15	63.36	64.56	3.01	255.18	221.1	44.33
2056–2065	18.97	0.074	63.83	65.05	66.27	3.22	268.03	227.0	46.13
2066–2075	18.91	0.076	63.68	64.90	66.12	3.25	262.91	226.9	46.49
2076–2085	19.00	0.079	63.85	65.08	66.31	3.33	265.22	228.1	47.29
2086–2095	19.03	0.077	63.89	65.11	66.34	3.29	265.59	227.1	46.94
2096–2105	18.18	0.075	62.41	63.60	64.79	3.23	260.32	224.2	46.35
2106–2115	17.87	0.069	60.89	62.06	63.24	3.05	256.06	219.7	44.67
2116–2125	17.23	0.053	57.98	59.14	60.31	2.57	233.14	206.9	39.64

Table 8: *Expected earthquake damage to Mexico City by decade, billions of constant 2026 MXN, undiscounted and not adjusted for future inflation. LCI/UCI are the 95% confidence interval on the expected value from Monte Carlo error ($N = 2 \times 10^5$ centuries per bound); they are not predictive intervals. Mean and median are both given because they differ by a factor of about 21 at decade resolution: the loss distribution is dominated by rare large events, so the mean is not a typical decade. A median decade costs 5 bn MXN on the central bound, while the mean is 23 bn MXN; the 95th percentile runs 80–95 bn. For scale, the observed 2017 event cost 16 bn MXN. The low–high span reflects vulnerability-parameter uncertainty and dominates all other uncertainty by two orders of magnitude. The apparent rise then fall across decades is not a trend: it is the ETAS process relaxing from the historical catalogue’s conditioning toward its stationary rate, and the variation is well within the predictive spread.*

Discounting, if wanted for a budgeting application, changes the totals substantially. At

Year	Low		LCI	Central				High	
	mean	med.		mean	UCI	med.	p95	mean	med.
2026	1.383	<0.001	4.44	4.76	5.09	<0.001	2.64	16.9	0.003
2027	1.375	<0.001	4.59	4.92	5.24	<0.001	3.18	17.6	0.004
2028	1.537	<0.001	4.96	5.31	5.66	<0.001	3.44	18.7	0.006
2030	1.680	<0.001	5.42	5.79	6.15	<0.001	3.92	20.0	0.009
2035	1.739	<0.001	5.54	5.91	6.28	<0.001	4.20	20.6	0.012
2040	1.607	<0.001	5.41	5.76	6.10	<0.001	4.80	20.8	0.015
2045	1.943	<0.001	6.15	6.54	6.94	<0.001	4.80	22.2	0.018
2050	1.910	<0.001	5.94	6.33	6.72	<0.001	5.32	22.1	0.020
2055	1.928	<0.001	6.20	6.59	6.99	<0.001	5.40	22.8	0.021
2065	1.928	<0.001	6.23	6.63	7.02	<0.001	5.56	22.9	0.022
2075	1.909	<0.001	6.26	6.65	7.04	<0.001	5.51	23.1	0.024
2085	1.805	<0.001	6.03	6.41	6.78	<0.001	5.67	22.7	0.024
2095	1.774	<0.001	5.88	6.25	6.62	<0.001	5.50	22.3	0.023
2105	1.812	<0.001	6.02	6.40	6.77	<0.001	5.35	22.4	0.020
2115	1.776	<0.001	5.67	6.05	6.42	<0.001	4.87	21.3	0.018
2125	1.519	<0.001	5.00	5.34	5.69	<0.001	3.64	19.0	0.007

Table 9: Expected annual damage for selected years, billions of constant 2026 MXN. Selected years are shown; the complete 100-year series is in `valor_presente_por_anio.csv`. At annual resolution the mean exceeds the median by more than five orders of magnitude on the central bound, because in a typical year nothing damaging happens: 21.5–32.2% of simulated years have exactly zero loss and the median year costs under 0.02 bn MXN. The mean is carried almost entirely by the rare years containing a large event, which is why the p95 column is the more useful planning figure at this resolution. The gap closes as the horizon lengthens: mean/median is $\approx 2 \times 10^5$ annually, 21 per decade, and 2.15 over the full century, as aggregation makes a damaging event increasingly likely to have occurred.

a 3% real discount rate the central 100-year present value falls from 624 bn MXN to 198.5 bn MXN, and at 5% to 124.5 bn MXN; the corresponding low and high bounds are 4.4–858 and 2.8–538 bn respectively. The full annual series and all three discount rates are in `valor_presente_serie_anual.csv` and `valor_presente_acumulado.csv`.

7.3 Computational cost

The full experiment is 600,000 independent century-long realisations (200,000 per soil bound), spanning 6×10^7 simulated years and approximately 2.19×10^9 simulated earthquakes, evaluated against an exposure base of 2,295,465 cadastral units aggregated to 3,472 one-kilometre cells.

7.4 Known limitations

1. **Peak ground acceleration is the wrong intensity measure for this city.** Damage in both 1985 and 2017 concentrated in mid-rise buildings whose natural period matched the site period of the lake-bed sediments [6]. A PGA-based model cannot express that resonance and therefore misallocates damage across building heights, even when the total is calibrated correctly. A spectral treatment requires building-period information the cadastre does not contain.
2. **The 100-year horizon assumes a frozen city.** Building stock, construction quality, and cadastral values are held at their 2021 values for the whole projection. All three will change, and the post-2004 share of the stock will grow, which biases the projection toward pessimism.
3. **The attenuation relation is anchored, not regressed.** Only the intercepts are fitted, to two observations; the slopes are literature-typical values. The out-of-sample test at 739 km under-predicts threefold. In future works we aim to generalize this idea.
4. **Spatial and temporal clustering are modelled separately.** ETAS supplies times and magnitudes; locations are resampled independently from the historical epicentre field. Aftershocks are therefore not placed near their mainshocks. For Mexico City loss this matters little, since all sources are distant and the distance distribution dominates, but it would matter for a local hazard question.

7.5 Next steps

In order of expected value: transcribe the *Valores Unitarios de las Construcciones* tables for 2021 and compute full cadastral value, land plus construction, in place of the assumed replacement cost; obtain the NTC-Sismo microzonation and replace the Atlas proxy; substitute published attenuation coefficients [1, 2, 3]; obtain an independent estimate of reporting completeness from insurance-claim or reconstruction-programme records, which would fix ρ externally rather than by fit and is the single most valuable addition to the calibration; add a second calibration event with a contemporaneous

593 exposure reconstruction, which the 2021 register cannot supply for 1985; move from
594 PGA to spectral acceleration with a building-height proxy; and introduce insurance
595 penetration data to convert gross loss into the insured loss relevant to the sector.

596 Reproducibility

597 All results are produced by the `pysismos` package accompanying this paper. The
598 pipeline runs in stages (catalogue cleaning, clustering, Voronoi weighting, ETAS fitting,
599 exposure construction, calibration, Monte Carlo), each independently re-runnable, with
600 parameters in a single configuration object and row-level accounting written for every
601 filtering step.

602 Acknowledgements and tooling

603 The analysis pipeline, statistical calibration, Monte Carlo engine and figures in this paper
604 were developed with Claude (Anthropic), used as an interactive analysis environment.
605 Its contributions included: porting the original R scripts to a staged Python package
606 and identifying seven data-preparation defects in the process; establishing the 1980
607 magnitude of completeness from b -value stability; designing and implementing the four-
608 observable joint likelihood calibration, including the reporting-completeness parameter
609 that reconciles a physical damage model with a reported-loss total; deriving the exact
610 magnitude-shift reduction that makes large-scale simulation feasible; and producing all
611 figures. The author directed the research questions, supplied the domain knowledge
612 and the calibration data, corrected the source clustering, and is responsible for all
613 interpretations and for any remaining errors.

614 8 References

- 615 [1] Arroyo, D., García, D., Ordaz, M., Mora, M. A., & Singh, S. K. (2010). *Strong*
616 *ground-motion relations for Mexican interplate earthquakes*. Journal of Seismology,
617 14(4), 769–785.
- 618 [2] García, D., Singh, S. K., Herráiz, M., Ordaz, M., & Pacheco, J. F. (2005). *Inslab*
619 *earthquakes of central Mexico: Peak ground-motion parameters and response spectra*.
620 Bulletin of the Seismological Society of America, 95(6), 2272–2282.
- 621 [3] Jaimes, M. A., & García-Soto, A. D. (2020). *Updated ground motion prediction*
622 *model for Mexican intermediate-depth intraslab earthquakes including V/H ratios*.
623 Earthquake Spectra, 36(3), 1315–1338.
- 624 [4] Abrahamson, N., Gregor, N., & Addo, K. (2016). *BC Hydro ground motion*
625 *prediction equations for subduction earthquakes*. Earthquake Spectra, 32(1), 23–44.

- [5] Stewart, J. P., Douglas, J., Javanbarg, M., et al. (2015). *Selection of ground motion prediction equations for the Global Earthquake Model*. Earthquake Spectra, 31(1), 19–45.
- [6] Singh, S. K., Reinoso, E., Arroyo, D., Ordaz, M., Cruz-Atienza, V. M., Pérez-Campos, X., & Iglesias, A. (2018). *Deadly intraslab Mexico earthquake of 19 September 2017 (M_w 7.1): Ground motion and damage pattern in Mexico City*. Seismological Research Letters, 89(6), 2193–2203.
- [7] Ordaz, M., Singh, S. K., Reinoso, E., Lermo, J., Espinosa, J. M., & Domínguez, T. (2018). *The 19 September 2017 M 7.1 Puebla–Morelos earthquake: Spectral ratios confirm Mexico City zoning*. Bulletin of the Seismological Society of America.
- [8] Galvis, F., Miranda, E., Heresi, P., Dávalos, H., & Silos, J. R. (2017). *Preliminary statistics of collapsed buildings in Mexico City in the September 19, 2017 Puebla–Morelos earthquake*. John A. Blume Earthquake Engineering Center, Stanford University.
- [9] Galvis, F. A., Miranda, E., Heresi, P., Dávalos, H., & Silos, J. R. (2020). *Overview of collapsed buildings in Mexico City after the 19 September 2017 (M_w 7.1) earthquake*. Earthquake Spectra, 36(2_suppl), 83–109.
- [10] Centro Nacional de Prevención de Desastres (2018). *Impacto socioeconómico de los desastres en México durante 2017*. Secretaría de Gobernación, México.
- [11] Gobierno de la Ciudad de México (2018). *Programa para la Reconstrucción, Recuperación y Transformación de la Ciudad de México*. Comisión para la Reconstrucción, CDMX.
- [12] Gobierno de la Ciudad de México (2017). *Atlas de Riesgos de la Ciudad de México: zona cero 2017*. Geospatial dataset, Portal de Datos Abiertos CDMX.
- [13] Bray, A., Wong, K., Barr, C. D., & Schoenberg, F. P. (2014). *Voronoi residual analysis of spatial point process models with applications to California earthquake forecasts*. Annals of Applied Statistics, 8(4), 2247–2267.
- [14] Bairakdar, R., Dupuis, D., & Mailhot, M. (2024). *Deviance Voronoi residuals for space-time point process models: An application to earthquake insurance risk*. arXiv preprint arXiv:2410.04369.
- [15] Omori, F. (1894). *On the aftershocks of earthquakes*. Journal of the College of Science, Imperial University of Tokyo, 7, 111–200.
- [16] Utsu, T. (1961). *A statistical study on the occurrence of aftershocks*. Geophysical Magazine, 30, 521–605.
- [17] Utsu, T., Ogata, Y., & Matsu’ura, R. S. (1995). *The centenary of the Omori formula for a decay law of aftershock activity*. Journal of Physics of the Earth, 43(1), 1–33. DOI: 10.4294/jpe1952.43.1

- 663 [18] Ogata, Y. (1988). *Statistical models for earthquake occurrences and residual analysis*
664 *for point processes*. Journal of the American Statistical Association, 83(401), 9–27.
- 665 [19] Aki, K. (1965). *Maximum likelihood estimate of b in the formula $\log N = a - bM$ and*
666 *its confidence limits*. Bulletin of the Earthquake Research Institute, 43, 237–239.
- 667 [20] Lee, I., & Lee, K. (2007). *Higher order Voronoi diagrams for disaster and emer-*
668 *gency management*. ISSNIP/ISPRS Joint International Workshop on Distributed
669 Geoinformatics.
- 670 [21] Gomez-Zapata, J. C., Brinckmann, N., Harig, S., et al. (2021). *Variable-resolution*
671 *building exposure modelling for earthquake and tsunami scenario-based risk assess-*
672 *ment: An application case in Lima, Peru*. Natural Hazards and Earth System
673 Sciences.