Beyond the floods: the economic case for adaptation and insurance in Brazil

Technical Annex[1]

Exposure layer

The exposure layer represents the spatial distribution of installed capital at risk in each city. For each of the eight municipalities, exposure is mapped on a grid of approximately 30 arcsec resolution (roughly one kilometre at the equator) using the LitPop methodology (Eberenz et al., 2020), in a two-stage allocation.

In the first stage, a national total of installed capital is allocated across the eight municipalities in proportion to each city's share of national GDP. The national total is the World Bank's produced capital series (Changing Wealth of Nations; World Bank, 2024), reported in real chained 2019 US dollars and scaled to the 2022 reference year using the growth of the World Bank's World Development Indicators GDP series. Each city's share is taken from the IBGE municipal accounts (Produto Interno Bruto dos Municípios), using 2022 values (pooled from Ipeadata, 2025). This stage fixes each city's aggregate exposure at its observed weight in the national economy.

In the second stage, each city's allocated total is distributed across grid cells using LitPop, whose weight for each cell is the product of night-time light intensity and population count. Cells that are brightly lit and densely populated receive a larger share of the city total than dark or sparsely populated cells, approximating the joint distribution of economic activity and settlement that drives where capital accumulates within the city. The two within-city inputs are:

- Night-time lights: the VIIRS night-time Day/Night annual band composite for 2022, accessed via Google Earth Engine (Elvidge et al., 2021).
- Gridded population: the IBGE's (2025) Statistical Grid based on the 2022 Population Census.

Three features of this approach matter for how the city-level estimates should be read. First, the within-city distribution rests on a satellite proxy rather than a direct measurement of where assets sit: the city totals are anchored to GDP, but the placement of capital within each city follows the lit-population surface, which Eberenz et al. (2020) note may not reproduce the spatial distribution of assets within metropolitan areas in all cases. Second, the relationship between nightlights and capital is weaker in informal settlements than in formal commercial districts, which can under-represent capital in rapidly expanding informal areas and over-represent it in long-established centres; this is most relevant for Manaus and Belém, where informal urban expansion has been significant. Third, the exposure layer is a 2022 snapshot held fixed over the 2025–2100 simulation period, so it reflects the 2022 distribution and level of capital rather than its future growth.

Holding exposure fixed at its 2022 level is a deliberate choice, not only a data constraint. Projecting how each city's capital stock would grow and shift over the century would require strong assumptions

[1] Annex to: Lara Miranda I, Martínez JP, Yusuf A, Gimenes F, Gonçalves E and Surminski S (2026) *Beyond the floods: the economic case for adaptation and insurance in Brazil.* London: Centre for Economic Transition Expertise, London School of Economics and Political Science. Full report available at: https://cetex.org/publications/beyond-the-floods-the-economic-case-for-adaptation-and-insurance-in-brazil

about future urban and economic development that the available data cannot support. It would also require a weighting scheme that differentiates losses in later years against earlier ones. The analysis stays neutral on this point: every year in the simulation window is valued equally, with no discounting applied, so the reported figures isolate the effect of the changing flood hazard on a constant 2022 asset base. One consequence is that the estimates do not capture the additional assets that urban growth is likely to place in flood-prone areas over the century, so on this dimension they are conservative in terms of future exposure. These features could be relaxed in future work with richer spatial data, such as building-footprint datasets, municipal cadastral records or built-up area indices that distinguish formal and informal urban fabric, as discussed in the limitations subsection below.

Hazard layers and model exclusions

The hazard layer captures the depth, frequency and spatial pattern of flooding across each city under different climate trajectories. The analysis draws on the ISIMIP2b ensemble of fluvial flood simulations (Zimmer et al., 2023), which provides a coordinated set of hazard projections suitable for cross-city comparison.

The ISIMIP2b protocol

The Inter-Sectoral Impact Model Intercomparison Project (ISIMIP) is an international initiative that coordinates climate impact modelling under standardised inputs. The 2b protocol provides fluvial flood inundation simulations generated by the CaMa-Flood global river-routing model (Yamazaki et al., 2011), driven by an ensemble of global hydrological models (GHMs) and global climate models (GCMs).

For each city, the analysis uses simulated annual maximum flood depth at a resolution of 150 arcsec (roughly five kilometres at the equator) over the period 2025–2100. The 76-year window provides 76 annual events per simulation. The full ensemble consists of six GHMs and four GCMs, giving up to 24 GHM × GCM combinations per city, repeated under three Representative Concentration Pathways (RCP 2.6, RCP 6.0 and RCP 8.5):

- GHMs: CLM45, CWatM, H08, LPJmL, MATSIRO, WaterGAP2
- GCMs: GFDL-ESM2M, HadGEM2-ES, IPSL-CM5A-LR, MIROC5

Within each retained model combination, average annual impact is computed under three flood protection assumptions (no protection, FLOPROS, and a 100-year return period), using installed capital as the exposure measure. Unless otherwise stated, the main text reports results under FLOPROS protection with installed capital exposure.

Looking for anomalous model combinations

When we plotted the FLOPROS-baseline average annual loss across the full ensemble for each city, a handful of combinations stood out as sitting well above the rest of the ensemble (city, GHM, GCM), often by a factor of 5-10, consistently across the three RCPs. This ensemble is an ensemble of opportunity: all models are not built to be mutually comparable, so a high dispersion across models is not necessarily a reason to treat any combination as anomalous. High values are not immediately excluded. Instead, a persistent, scenario-independent pattern allows the selection of combinations for closer inspection. Specific combinations were excluded only where their behaviour could be traced to a specific, independently documented model weakness. Combinations that were merely high, with no such documented explanation, were retained.

Figure A.1 is the diagnostic employed to surface these patterns. Each panel shows average annual loss for one city under one RCP, with rows for the six GHMs and columns for the four GCMs. Cell colour shows how far each combination sits from the city-RCP ensemble median, on a log scale; cell numbers show the absolute value in US$ million per year. The patterns of interest are those where a row (one GHM driving anomalous results across most GCMs) or a column (one GCM driving anomalous results across most GHMs) lights up red across multiple RCPs for the same city.

Figure A.1. Heatmap of average annual loss by GHM × GCM × RCP for each city

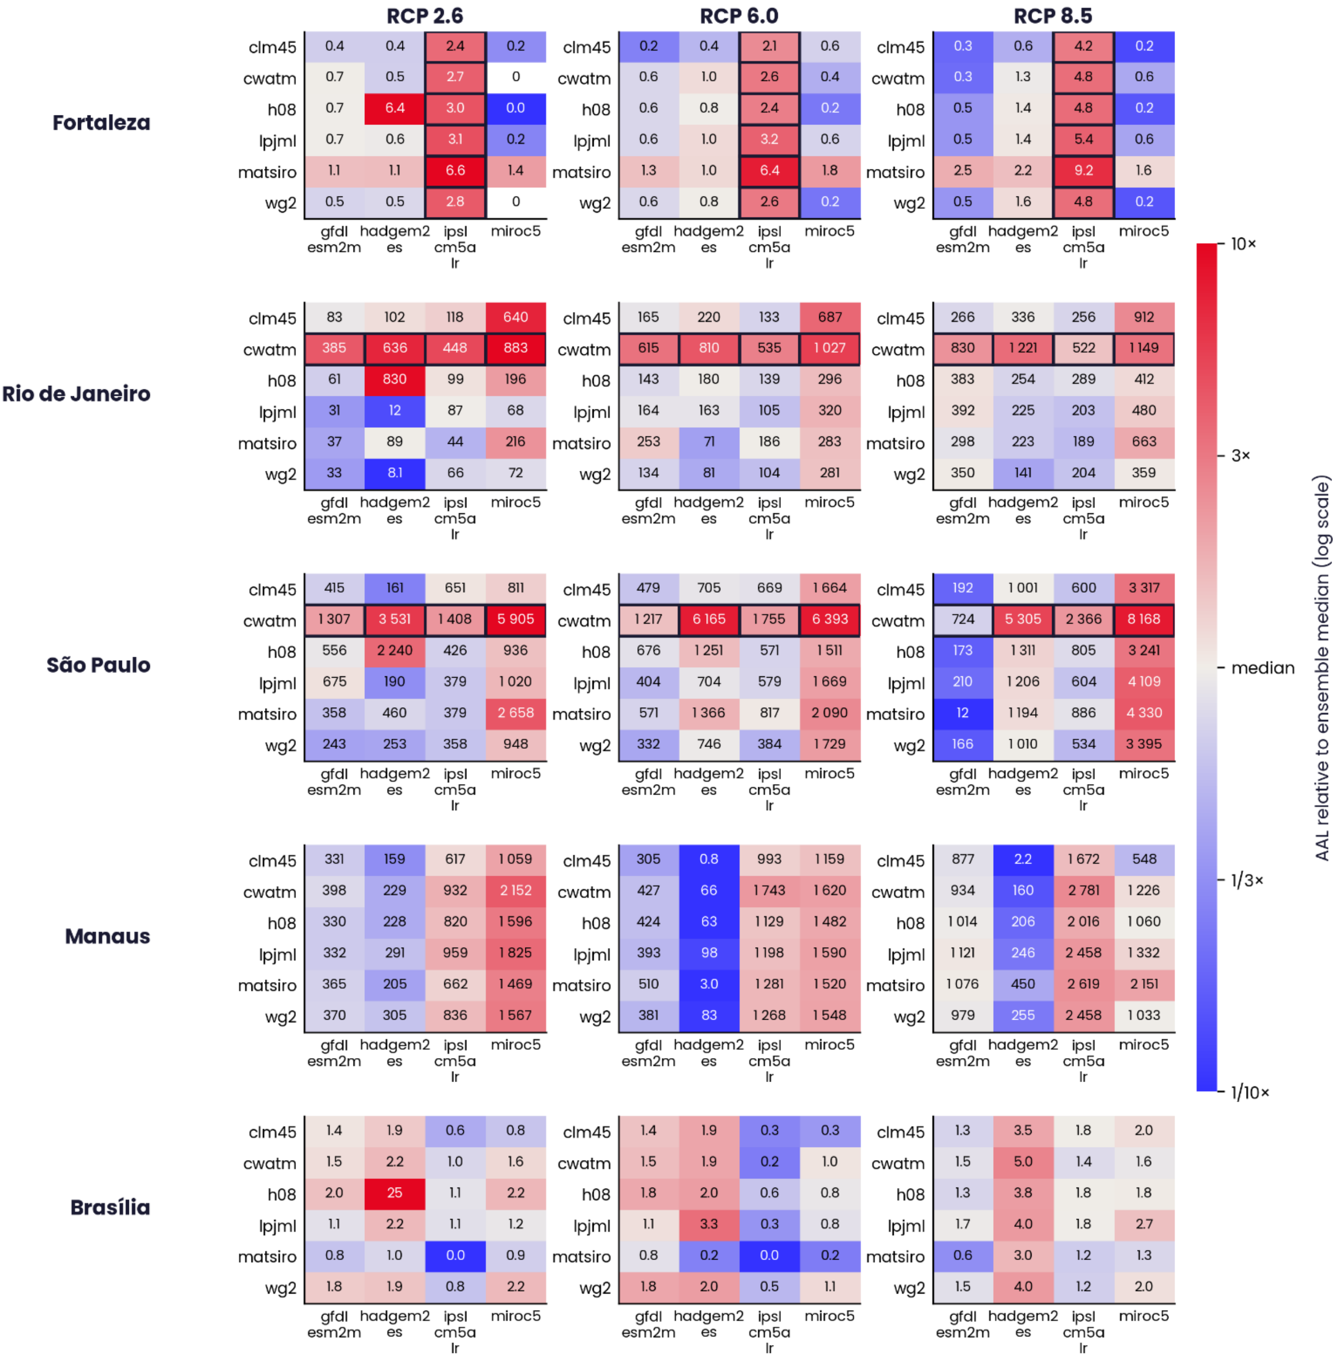

Note: Model-selection heatmap for three cities with documented exclusions (Fortaleza, Rio de Janeiro, São Paulo) and two control cities (Manaus, Brasília). Cells show average annual loss (AAL) in US$ millions per year for each hydrological model (GHM, rows) × climate model (GCM, columns) combination at FLOPROS protection. Colour indicates the ratio of the cell's AAL to the city × RCP ensemble median on a log scale. Dark-blue borders mark cells excluded by the Part 1 procedure. Note: wg2: watergap2.

Source: Authors' calculations using Zimmer et al. (2023) and Bresch and Aznar-Siguan (2021).

Which combinations were excluded

Each candidate flagged by the heatmap diagnostic was then checked against the published model-evaluation literature. The question was whether the anomaly could be plausibly explained by a documented model weakness, in which case excluding it would tighten the ensemble. Alternatively, it might reflect a legitimate extreme physical response that should be retained. The resulting exclusions affected 14 of the 192 city × GHM × GCM combinations in the full grid (around 7%).

For Fortaleza, all combinations driven by IPSL-CM5A-LR were dropped. Under this GCM, all six GHMs produced losses four to ten times the ensemble median for Fortaleza, while behaving normally under the other three GCMs. The fact that the anomaly is the same across all GHMs but specific to one GCM points to the precipitation forcing rather than the runoff scheme. Beck et al. (2017) suggest that precipitation biases can be amplified into much larger runoff biases in semi-arid regions, which fits Fortaleza's Caatinga setting.

For Rio de Janeiro and São Paulo, all CWatM combinations were dropped. CWatM produced elevated losses across three of the four GCMs (factors of 4 to 10× the ensemble median), while the other five GHMs behaved normally under the same forcings. The pattern is GHM-specific rather than GCM-specific, which would point to how CWatM generates runoff rather than in any climate forcing. CWatM in ISIMIP2b uses an uncalibrated global default parameter set rather than basin-specific calibration, and our re-validation confirms that this elevated-loss pattern for both cities is stable under the updated exposure layer. Flaws in the simulated discharge of uncalibrated global hydrological models of this kind have been documented in the literature (Burek et al., 2020; Zaherpour et al., 2018), which further supports the decision to exclude these combinations. The overall effect of the exclusions on the reported losses is shown in the following subsection (Figure A.2).

Combinations that could have been excluded but were kept

- LPJmL for Manaus and Belém. LPJmL is the only GHM in the ensemble with dynamic vegetation (Schaphoff et al., 2018). For cities where deforestation can plausibly influence the runoff regime (Manaus in the central Amazon, Belém at the Amazon mouth) this is an important channel to be captured. Removing LPJmL would mean dropping the only ensemble member that responds to land use change in cities where land-use change matters.

- MATSIRO + MIROC5 for São Paulo. Flagged in the heatmap at around 4.5× the ensemble median, but retained because Zaherpour et al. (2018) identify MATSIRO as one of the few GHMs in the ensemble that does *not* tend to overestimate high flows.

- IPSL-CM5A-LR for Belém under RCP 8.5. Elevated by a factor of around 4× across all GHMs but only at RCP 8.5. The single-RCP nature of the pattern would be consistent with a real high-emissions precipitation response rather than the cross-RCP pattern that would suggest a model defect.

Losses without exclusions

Figure A.2 shows the effect of the exclusions on the reported results. For each city it compares the full ensemble, before any combination is removed, with the reported ensemble used in the main text, and marks the excluded combinations individually. In the three cities with exclusions (Fortaleza, Rio de Janeiro and São Paulo) the excluded combinations all fall at the high-loss end of the distribution: across these cities they run roughly four to five times the median of the retained combinations and sit around the 85th to 90th percentile of the full ensemble. Removing them therefore lowers the reported central estimate for each affected city. For the remaining five cities the exclusions make no difference, since no combination was removed. In this specific sense the reported estimates are more conservative than the full ensemble: the exclusions strip out the combinations that would have raised the headline losses, not lowered them.

Figure A.2. Average annual loss as a share of city GDP, full ensemble versus reported estimates

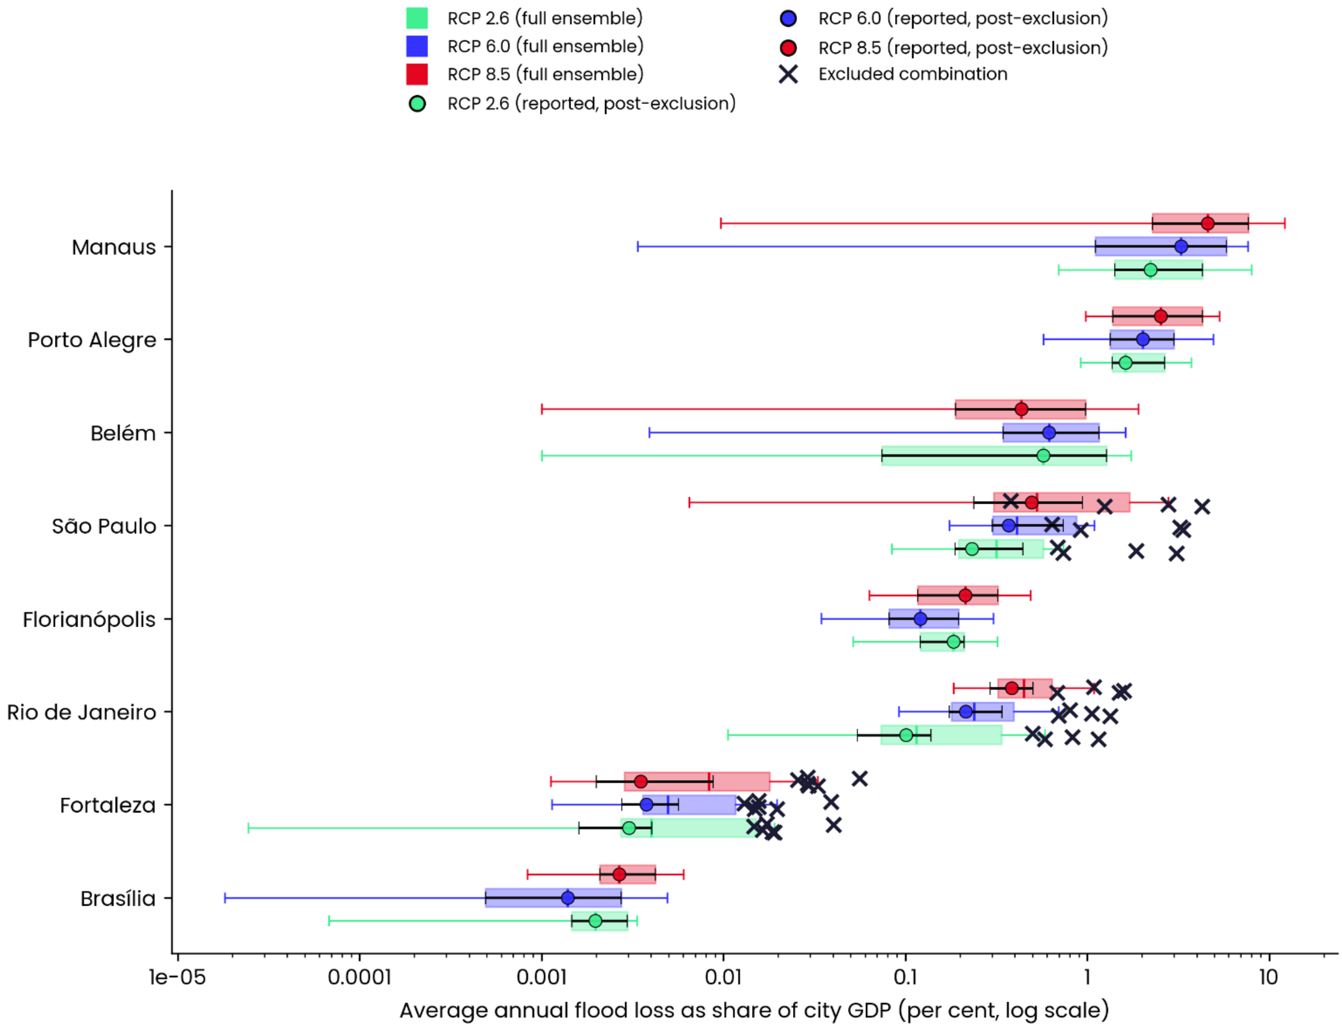

Source: Authors' calculations using Zimmer et al. (2023) and Bresch and Aznar-Siguan (2021).

Impact computation

Hazard and exposure are combined to produce monetary loss estimates through the CLIMADA platform (Bresch and Aznar-Siguan, 2021). The translation from flood depth to damage relies on a depth-damage function, which specifies the fraction of asset value destroyed at any given flood depth.

The depth-damage function

The analysis uses the JRC RF6 depth-damage function for South America (Huizinga et al., 2017), which specifies two quantities at a set of intensity knot points: the mean damage degree (the fraction of value destroyed at a given depth, conditional on being affected) and the percentage of affected assets (the fraction of assets in a grid cell that suffer damage at a given depth). In the RF6 specification, damage rises monotonically from zero at the depth-onset threshold to full destruction at depths above approximately six metres.

Average annual loss

For each event in the simulated record, total monetary loss at the city level is computed by applying the depth-damage function to each grid cell, multiplying by the produced capital exposure in that cell, and summing across cells. This produces a loss value for each of the 76 simulated years in the ensemble.

The average annual loss (AAL) is then the simple mean of the per-event losses across the 76 years in the record $T = 76$:

$$AAL = \frac{1}{T}\sum_{T} L_t$$

where L_t is the total city loss in simulated year t. Every year in the record receives equal weight, on the assumption that each simulated year represents an independent draw from the underlying flood-depth distribution. The average annual loss can also be interpreted as the area under the city's annual loss-exceedance probability curve over the simulated record, though that interpretation is constrained by the same 76-year sample.

Conversion to share of GDP

Average annual loss is reported in constant 2019 US dollars throughout. To express losses as a share of city GDP, the denominator is the same municipal GDP used to allocate exposure: each city's share of national GDP (IBGE municipal accounts, 2022 values; Ipeadata, 2025) applied to Brazil's GDP in constant 2019 US dollars. Because the numerator and denominator share a constant 2019 US dollar basis and the same 2022 reference year, the average annual loss to GDP ratios are consistent in both price base and timing. On this basis, city GDP ranges from approximately US$5.3 billion for Florianópolis to US$190.5 billion for São Paulo.

Losses through time

The average annual loss reported throughout this analysis is a single figure for each city, scenario and protection standard, taken as the mean over the full 2025–2100 window. It therefore treats the flood hazard as if its average intensity were constant across the century, whereas in practice the hazard is expected to evolve as emissions accumulate. Figure A.3 relaxes this assumption by disaggregating the FLOPROS-standard losses into decade-level distributions for three illustrative cities (Manaus, Porto Alegre and Florianópolis). The plot also includes the estimates previously presented in Figure 4.1 for comparison.

Two features stand out. First, for several city–scenario combinations the decade-median loss tends to rise over the century, with the later decades sitting above the earlier ones. This is most apparent under the higher-emissions pathway and is consistent with flood damage becoming more acute as cumulative emissions, and hence warming, increase. The separation between the lower- and higher-emissions pathways is correspondingly somewhat clearer in the later decades, although the wide decade-level uncertainty makes it difficult to distinguish cleanly.

Second, disaggregating the record into decades sharply widens the spread within each period: with far fewer years contributing to each estimate, the decade-level ranges are much wider than the whole-period range, reflecting the deep uncertainty that arises from the modelling, the underlying data and the scenario design. Thus, given that the hazard layer is not constant throughout the century, the averaging across the entire 76 years will tend to overstate the losses of the earlier decades and understate those of the later ones, so it is best read as a century-scale summary rather than a description of any single decade.

Figure A.3. Flood losses as a share of city GDP by decade

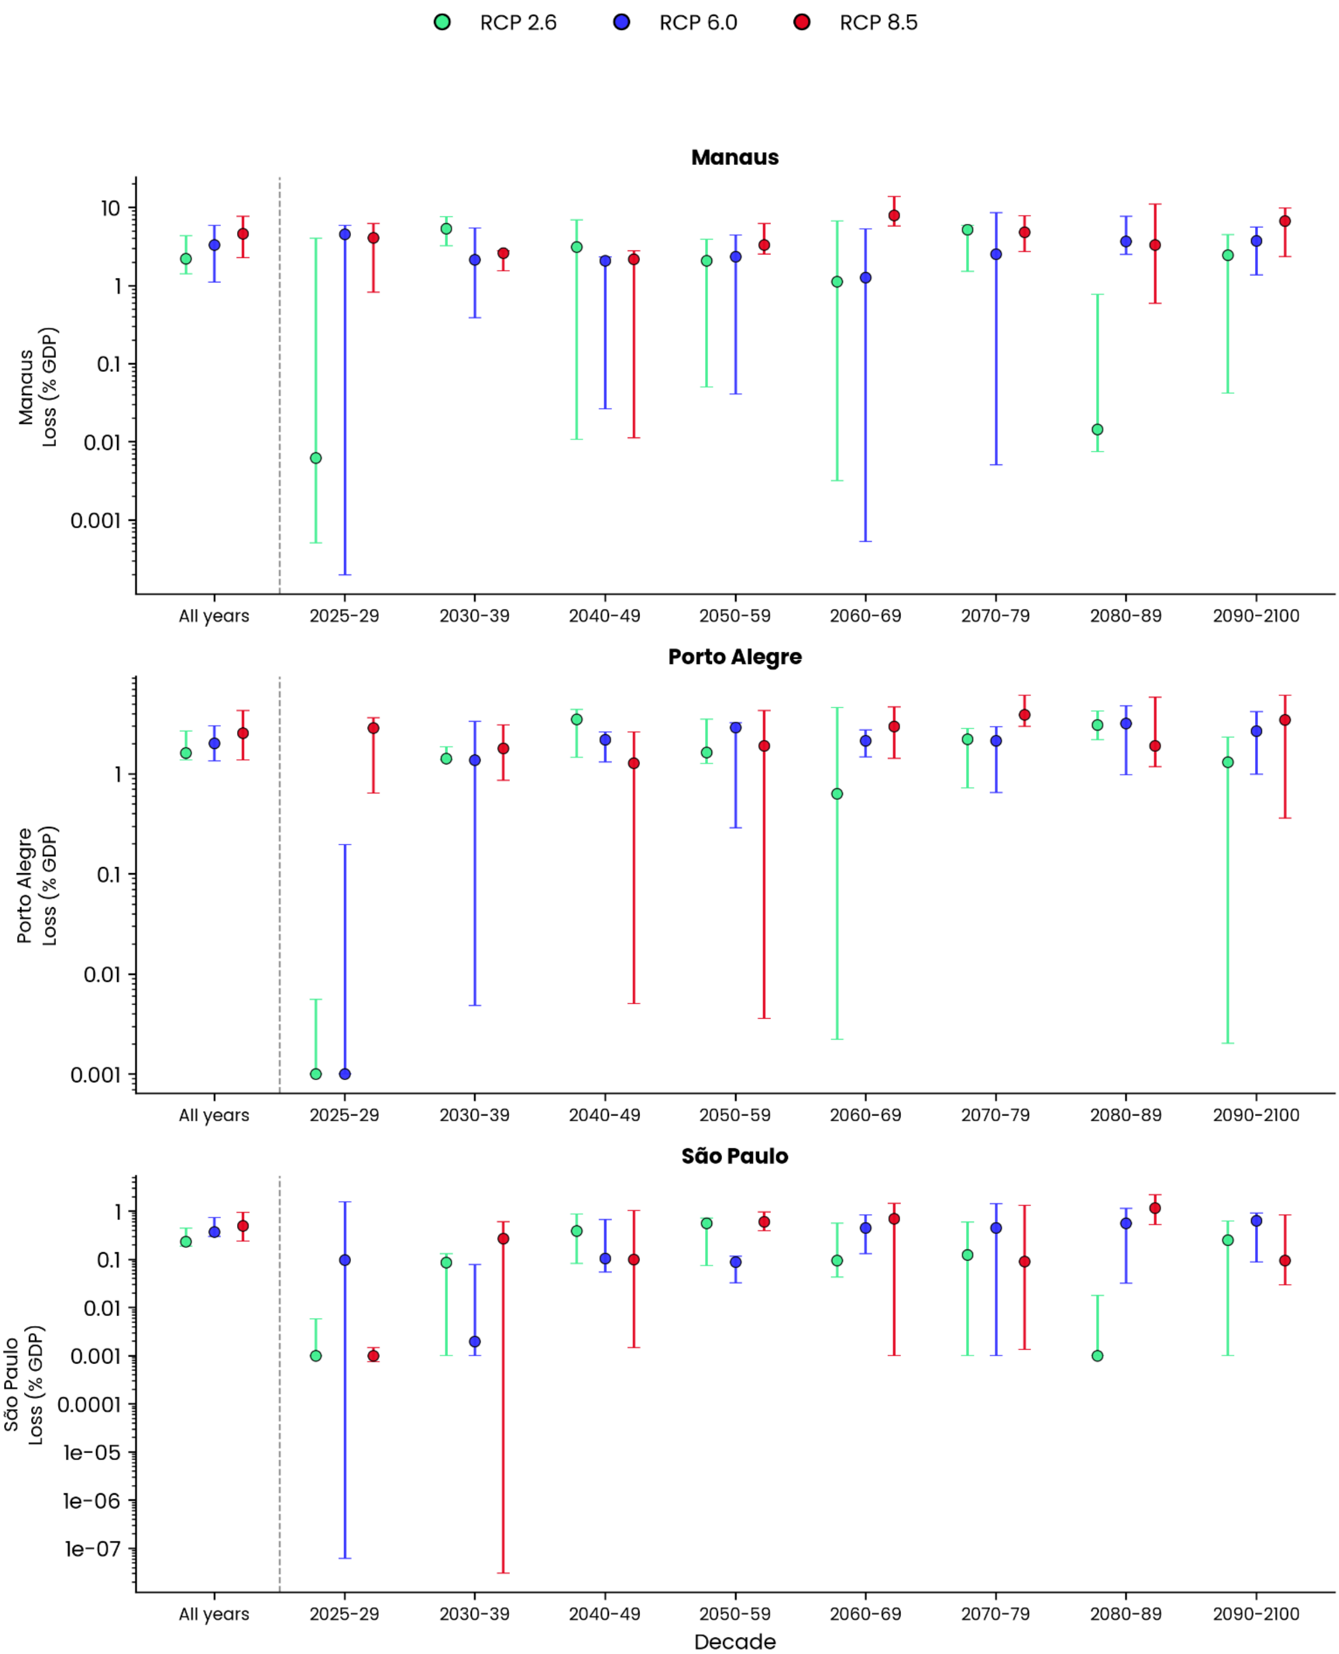

Note: Points are ensemble medians across the retained model combinations, and the vertical bars show the 25th and 75th percentiles. "All years" is the average across the entire century, consistent with Figure 4.2. Source: Authors' calculations using Zimmer et al. (2023) and Bresch and Aznar-Siguan (2021).

Estimating expected losses across protection standards

A central question for the analysis is how average annual flood losses change as cities upgrade their flood protection. Answering it requires a method for translating a notional protection standard, typically expressed as the return period of the design flood, into a modification of the underlying hazard layer. This subsection sets out the methodology used to produce the loss-protection curves shown in Figure 4.2.

The ideal approach and why it cannot be applied

The textbook way to apply a return-period-based protection standard is to work centroid by centroid: for each grid cell in the hazard layer, fit an exceedance curve to the 76 annual maximum flood depths simulated for that cell, find the depth that corresponds to the target return period, and zero out any flood event at that cell shallower than that depth.

Two features of the data make this approach unreliable for the cities studied here. The first is flood sparsity: most centroids in each city's bounding box flood in only a small number of the 76 simulated years, often just two or three. Fitting an exceedance curve to such a short and sparse record produces predicted protection depths that quickly exceed the maximum observed depth at that centroid, meaning the modification effectively zeroes the centroid entirely regardless of the protection level being modelled. The second is the record length itself: 76 years can only directly inform return periods up to around 76, and any inference about a 100-year event requires extrapolation beyond the observed range, which cannot be validated from the data alone.

The combination of these two features means that the per-centroid approach produces protection-standard estimates that are either degenerate or unverifiable. With longer simulation records or richer per-centroid distributional fitting techniques such as generalised extreme value modelling, the textbook approach would be viable. With the data currently available, an alternative is needed.

The depth-sweep approach

The alternative used in this analysis sweeps a single flood protection depth threshold d_p across the city, rather than computing a return-period-specific depth for each centroid. For any value of d_p, the depth-damage function is modified so that mean damage degree is set to zero for all flood depths below d_p, while the function is left unchanged at depths at or above d_p. The CLIMADA impact calculation is then re-run with the modified damage function, producing an average annual loss conditional on that protection level. Repeating the calculation across a fine grid of d_p values traces out a loss-protection curve for each model combination.

This approach trades the centroid-by-centroid spatial resolution for a city-wide protection abstraction. Three strong assumptions underpin it. First, the threshold d_p is uniform across the city: every centroid is assumed to be protected to the same flood depth, when flood protection infrastructure is spatially heterogeneous and tends to concentrate in areas with the highest historical loss exposure.

Second, the protection is treated as binary: floods shallower than d_p are fully prevented, and floods at or above d_p are fully unmitigated, with no attenuation in between. Real infrastructure typically reduces peak depths above the design standard rather than passing them through unchanged, but this effect is below the five kilometre resolution of the CaMa-Flood hazard layer and cannot be modelled here. Third, the assignment of a return period to each value of d_p, set out below, relies on an empirical translation that conflates spatial and temporal variability in the simulated flood record.

These may be strong assumptions, and the loss-protection curves should be read with them in mind. Nevertheless, they do not undermine the central policy message of the analysis. These curves are not designed to determine the optimal protection standard for any single city, which would require local engineering analysis well beyond the scope of this exercise. Their purpose is to show that stronger flood protection produces meaningful reductions in average annual losses, and that those reductions translate into reduced fiscal vulnerability in the most exposed cities. That conclusion is robust to the simplifications above. The case for adaptation investment, and for the parallel development of insurance solutions to absorb the residual risk that adaptation cannot eliminate, does not depend on the precise shape of any individual curve.

Depth-sweep implementation

For each city, the analysis covers a set of GHM × GCM × RCP combinations. The depth sweep is implemented as a discrete grid of protection threshold values. Each depth value d_p on the grid triggers one CLIMADA impact calculation with the corresponding modified damage function producing one average annual loss value. The resulting output is a staircase curve of average annual loss against protection depth, with the discrete jumps reflecting the structure of the underlying calculation and data. This reflects the discrete structure of the underlying calculation and data. At each new threshold d_p, the tranche of grid cells with flood depths between the previous and current threshold drops out of the damage calculation, producing a discrete step downward in average annual loss. Panel A of Figure A.4 shows the staircase shape resulting from the depth-sweep on a specific model combination in Porto Alegre.

From protection depth to return period

The depth-sweep approach produces a curve in protection-depth space, but the policy-relevant axis is the return period of the design flood. The translation from one to the other relies on an empirical cumulative distribution function (CDF) of non-zero flood depths in the simulated record. For each combination, all non-zero (event, centroid) depth values across the 76 simulated years are pooled, and the empirical exceedance frequency at any depth d_p is computed as the fraction of pooled values that exceed d_p. Formally, let $d_{c,t}$ be the maximum flood depth in centroid c in year t. The CDF can be expressed as:

$$\hat{F}(d_p) = \frac{1}{T}\frac{1}{C}\sum_{T}\sum_{C} \mathbb{I}[0 < d_{c,t} \leq d_p]$$

Where $\mathbb{I}$ is the indicator function, which enables the aggregation of all the non-zero flood depths below the threshold d_p. From this equation one can compute an implicit return period as:

$$RP(d_p) = \frac{1}{1 - \hat{F}(d_p)}$$

This translation rests on a substantive assumption: that pooled (event, centroid) depth pairs are exchangeable, meaning that depth variability across centroids each year is treated as comparable to depth variability across years for a given centroid. This assumption conflates spatial heterogeneity (different centroids flood at different depths in a single event) with temporal heterogeneity (return periods across the simulated record). The resulting return-period axis is therefore city-specific and combination-specific, and not directly comparable across cities. The translation is also capped at the record length: any value of d_p beyond the maximum observed depth in the simulation cannot be assigned a return period from the data alone.

Expressing the protection threshold as a return period also situates the loss-protection curve in relation to the exceedance probability curve used in catastrophe risk analysis. The two representations carry the same frequency-severity information in different forms. A conventional exceedance probability curve reports the loss from a single event of a given return period, whereas the loss-protection curve reports the average annual loss that remains once all events up to that return period are prevented. As the protection standard rises, only progressively rarer and deeper floods still cause damage, so the residual expected loss falls, tracing what is in effect an inverse view of the exceedance curve. Read at standard return periods such as 10, 20, 50 and 100 years, the curve conveys both the severity of rare events and the average annual loss avoided by protecting against them. Panel B of Figure A.4 shows the transformation to return period space.

Fitting a power law

The staircase shape is an artefact of the methodology and the coarseness of the underlying data, whereas the physical relationship between protection and expected losses is expected to be smooth. The next step is therefore to fit a smooth curve through the staircase. Different sources of information are exploited to fit a smooth curve that mimics this phenomenon. Firstly, the step edges from the staircase plot are ideal candidates, these correspond to points where average annual losses drop by at least 0.1% of the unprotected baseline. Secondly, ISIMIP2b includes three pre-computed protection layers:

- No protection. Return period of 1 by definition. The corresponding average annual loss is the unprotected baseline computed by CLIMADA without any modification of the damage function.

- FLOPROS. The average annual loss after applying the FLOPROS protection standard (Scussolini et al. 2016), a global database that documents existing flood protection levels by region. Its associated return period varies by location and is inferred per combination.

- 100-year. The average annual loss after applying a 100-year flood protection standard.

Panel C of Figure A.4 presents the chosen points for fitting the smoother curve, which can be formally expressed as:

$$AAL(RP) = A \cdot RP^{-\lambda} \leftrightarrow \log[AAL(RP)] = \log(A) - \lambda \log(RP)$$

The linearised version of the relationship is fit by fixing the intercept $\log(A)$ to the unprotected baseline; thus the only parameter estimated from the data is the decay exponent λ. This exponent has a direct interpretation: it is the elasticity of average annual loss with respect to the protection standard, so a 1% increase in the return period is associated with a λ% decrease in expected losses.

The fit is performed under the no-protection setting of the hazard layer, which recovers the implicit loss-protection relationship without conflating the methodology with the protection assumption being modelled. The FLOPROS return period for each combination is then derived by inverting the power-law fit at the observed FLOPROS-level average annual loss. Panel D of Figure A.4 overlays the fitted curve on the staircase, step edges and anchor points.

Figure A.4. Progressive build-up of the loss-protection curve fitting procedure for one example combination

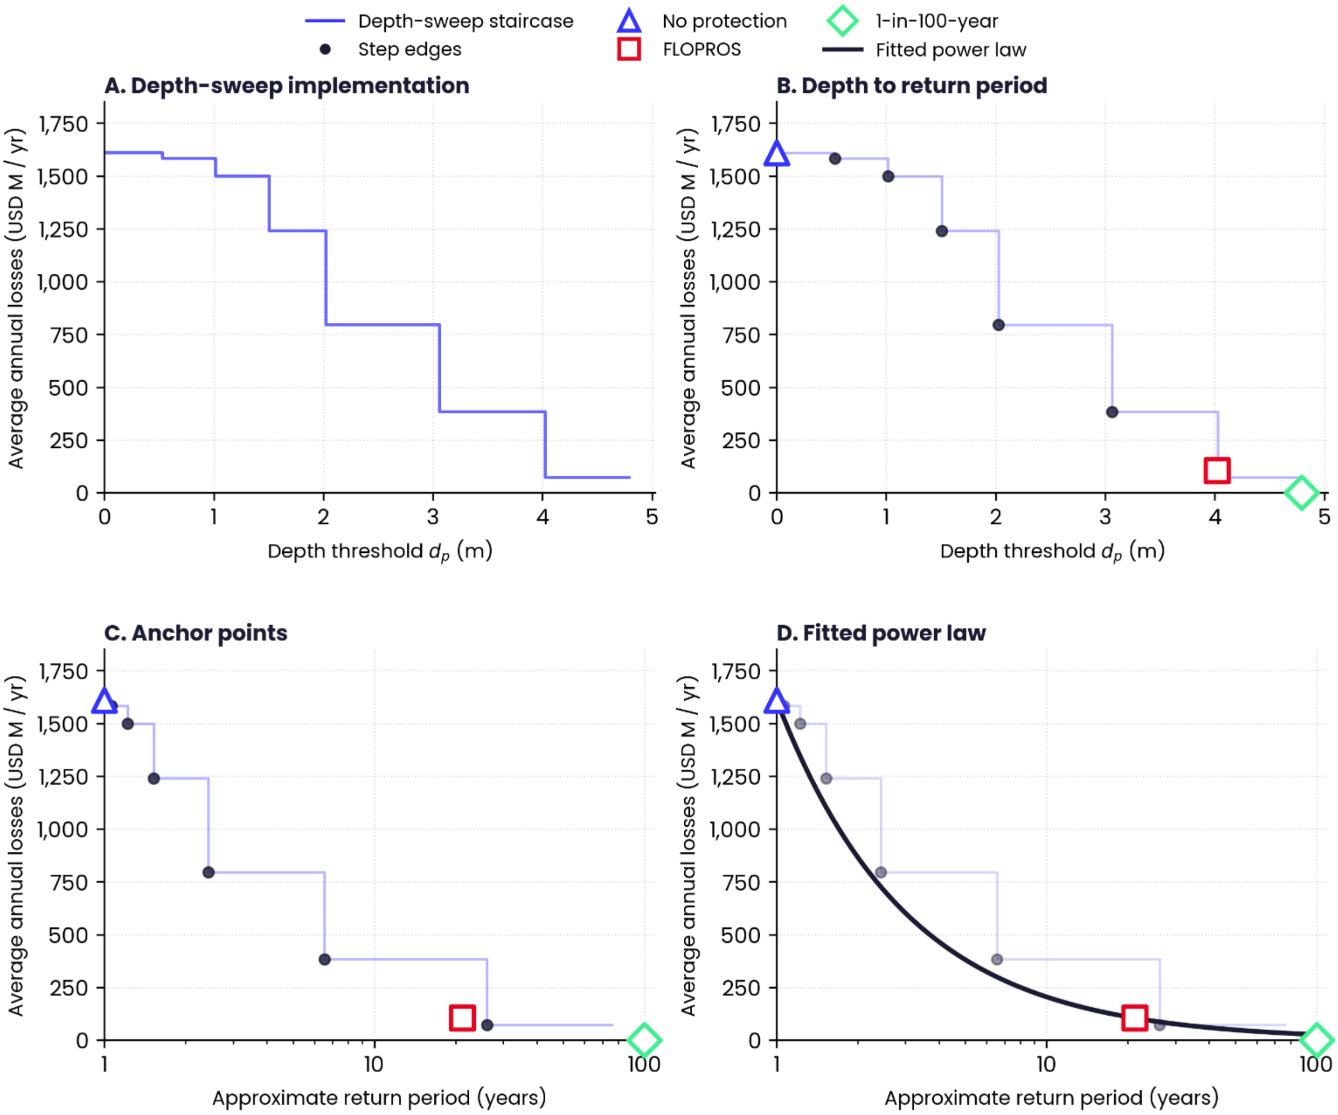

Note: Construction of the loss–protection curve, worked example for Porto Alegre under CLM4.5 × GFDL-ESM2M at RCP 6.0. Panel A: staircase of AAL against the depth threshold d_p from the CLIMADA-based depth sweep. Panel B: same staircase with the information-bearing step edges (dark-blue dots) and the three anchor points (No protection, FLOPROS, 100-year). Panel C: step edges and anchor points translated from the depth axis to the approximate return-period axis via the empirical CDF of non-zero flood depths. Panel D: fitted power law AAL(RP) = A · RP^(-λ) overlaid on Panel C, with A fixed to the No-protection anchor. The same procedure is applied to every retained (city, GHM, GCM, RCP) combination in the analysis.

Source: Authors' calculations using Zimmer et al. (2023) and Bresch and Aznar-Siguan (2021).

The fitted power law implies a constant loss elasticity across all protection levels, which may not hold for cities with a two-regime flood structure. Manaus is the clearest example: the Amazon floodplain produces a plateau in the depth-loss relationship between approximately 5 and 10 metres of flood depth, beyond which losses rise sharply again as central business district elevations are exceeded. A single power law cannot capture this structure exactly. In practice, the fit is good for most combinations (the coefficient of determination exceeds 0.95 for over 80% of fits) but is less reliable for combinations where the loss-depth relationship is genuinely two-regime.

Ensemble aggregation

The power law is fitted for each city and each GHM × GCM × RCP combination. Figure 4.2 presents all the curves on a common return-period grid running from 1 to 100 for each city. Specific quantiles of the curves' distribution are computed at each return period, namely, the median (dark blue), the inter-quartile range (25th to 75th percentiles, darker band) and the 5th to 95th percentile range (lighter band). The FLOPROS and 100-year anchor markers on Figure 4.2 are the medians of the per-combination average annual loss values at those two protection standards. For some cities the anchor markers do not sit exactly on the fitted median curve, nor the percentiles. This is because the curves smooth the depth-sweep staircases and the ISIMIP anchors into a single parametric form, while the anchor markers themselves reflect the raw ISIMIP impact calculations without smoothing.

Limitations and next steps

The methodology described in this appendix relies on several strong assumptions, each of which could be relaxed in future iterations of the analysis with richer data or more advanced techniques. This subsection draws together the methodological limitations flagged throughout the appendix and pairs each with a concrete suggestion for how it could be addressed.

Exposure

The exposure layer uses LitPop to distribute each city's installed capital across grid cells, with weights derived from 2022 nightlights and gridded population, once the city totals have been set by GDP share. This within-city step can under-represent capital in informal areas, and the 2022 inputs do not capture growth after that year. Moreover, LitPop has been validated against subnational GDP shares rather than against direct capital stock measurements. Three improvements that could help here are:

1. Building-footprint. Datasets such as Google Open Buildings or Microsoft Building Footprints provide a more direct measure of where built capital sits on the ground and would allow the disaggregation step to be re-weighted using built-up area rather than nightlights alone.

2. Municipal cadastral data. Where available from Brazilian state and city authorities, it would provide property-level valuations that could anchor the within-city distribution to local administrative records rather than to satellite proxies.

3. Dynamic exposure. This could be incorporated by projecting how produced capital is likely to grow over 2025–2100 in line with population and economic growth assumptions, rather than holding exposure fixed at the 2022 reference year. The current static exposure means that all scenario differences in the analysis reflect hazard changes only. Exposures that evolve would capture the additional fiscal liabilities that comes from urban growth in flood-prone areas.

Hazard

The hazard layer uses ISIMIP2b CaMa-Flood simulations at a resolution of 150 arcsec (around 5 km at the equator). This is sufficient to resolve large floodplains but coarse for the neighbourhood-level features that determine flood depth at the city scale. The analysis also covers fluvial flooding only. At roughly five kilometres, each flood-depth value covers an area spanning about 25 of the one kilometre exposure cells, so one depth is applied uniformly across ground whose elevation, and therefore true flood exposure, varies within the pixel. Because inundation is closely tied to local elevation, this can bias city-level losses in either direction. Therefore, the sign of the bias is hard to determine.

Three improvements would address these limitations:

1. Higher-resolution hazard layers are the most direct remedy for the resolution mismatch described above. Two routes are available: applying a downscaling algorithm that redistributes the coarse inundation using a high-resolution digital elevation model or adopting an already downscaled and yet more granular hazard product. Where available for Brazilian basins, these would allow the analysis to resolve flood depths at the neighbourhood scale rather than the basin scale, which could address the resolution-driven bias.

2. **Coastal storm surges** could be integrated by combining the fluvial hazard with output from the Global Tide and Surge Reanalysis (GTSR) or equivalent, which would be particularly relevant for Rio de Janeiro, Florianópolis and Belém.

3. **Pluvial flooding** from overwhelmed urban drainage is not modelled by global hazard products at all and would require coupling the analysis with city-level drainage capacity data or rainfall-runoff modelling. This is the largest hazard gap for São Paulo and Fortaleza.

Loss-protection curve methodology

The depth-sweep approach assumes a uniform protection threshold across each city, binary protection performance below and above that threshold, and an empirical translation from protection depth to return period that treats spatial and temporal variability as exchangeable. The power-law fit assumes a constant loss elasticity across all protection levels, which is less reliable for cities with two-regime flood structures.

The most impactful improvement would be longer simulation records or richer extreme value methods. Generalised extreme value (GEV) fitting at the per-centroid level, applied to a longer simulation record than the 76-year ISIMIP2b window, would make the textbook per-centroid approach viable and remove the need for the depth-sweep approximation entirely. Where ISIMIP3b or comparable next-generation simulation ensembles become available with longer records, this would be the natural upgrade. Spatially heterogeneous protection could be modelled by reading from FLOPROS or equivalent at the centroid level rather than imposing a single city-wide threshold.

Modelling uncertainty

The retained model ensemble covers six GHMs, four GCMs and three RCPs, with city-specific exclusions where the literature supports the diagnosis of a model defect. The ensemble captures a wide range of plausible futures, yet it cannot represent deep uncertainty. This uncertainty also has a temporal dimension: because the reported losses are averaged over the full 2025–2100 window, they do not capture how risk evolves within the century, which the decade-level disaggregation in Figure A.3 makes visible, both in the tendency for losses to grow towards the later decades and in the wider spread once the record is split into shorter periods. Ensemble expansion to incorporate additional GHM × GCM pairs from CMIP6-driven simulations would broaden the range. Probabilistic weighting of the ensemble members, rather than the equal weighting used here, would allow more credible weight to be given to models with documented skill in the relevant climate scenarios.

Taken together, these are the directions in which the analysis would naturally extend. These improvements would sharpen the precision of the city-specific estimates, broaden the range of hazards considered and tighten the confidence with which a given protection investment can be evaluated against its avoided-loss benefit.