

Probabilistic compound flood risk to critical infrastructure in a protected delta: loss concentration, sea-level rise, and tail-risk attribution

Victor Cárdenas-Santiago

victor.cardenas@ubc.ca

University of British Columbia

Research Article

Keywords: compound flooding, critical infrastructure, probabilistic risk assessment, loss exceedance curve, Fraser River delta, sea-level rise

Posted Date: October 8th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11065247/v1>

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

[Read Full License](#)

Additional Declarations: No competing interests reported.

Probabilistic compound flood risk to critical infrastructure in a protected delta: loss concentration, sea-level rise, and tail-risk attribution

Victor Cárdenas-Santiago

Institute for Resources, Environment and Sustainability, The University of British Columbia, Vancouver, Canada

Abstract

This paper presents a probabilistic compound coastal–fluvial flood risk assessment for 376 critical infrastructure assets of the Fraser River delta, Canada (CAD 55.61 billion replacement value), a census within five pre-declared categories. A 600,000-event stochastic catalog couples goodness-of-fit-tested bi-seasonal copulas to sampled tidal forcing and three sea-level tiers; a reduced-physics hydrodynamic model, evaluated at seven federal gauges with dikes represented explicitly, converts events to water surfaces; a Gaussian-process surrogate (cross-validated depth error 0.053 m) extends the physics to catalog scale; and a custom non-residential damage-function library prices losses. The configuration is Unsatisfactory under a published estuarine validation criterion and adopted under a fitness-for-purpose statement; five bias mechanisms share a conservative sign, so results are a lower bound with respect to them, with two-sided footprint-sampling and valuation uncertainty bracketed. Average annual loss (occurrence basis) rises from CAD 0.374 million today to 119.6, 350.6, and 1,726.2 million at +0.77, +1.00, and +1.50 m relative sea-level rise; the ratio of average annual loss to the 100-year probable maximum loss rises from 0.018 to 0.702; and losses concentrate on two wastewater plants; the regime transition survives within-site sampling at the dominant asset, absolute levels do not. An 1894-class scenario prices the provincial reference flood at CAD 4.8 million today against 2,030 million under the delta-adjusted tier. Statistical dependence between discharge and surge contributes at most ~ 3 percent of tail losses; hydrodynamic interaction is negligible for the median extreme event yet reaches ~ 30 percent in surge-driven events; compound tail risk is therefore governed primarily by marginal severities against protection freeboards.

Keywords: compound flooding, critical infrastructure, probabilistic risk assessment, loss exceedance curve, Fraser River delta, sea-level rise

1. Introduction

Climate change is converging on the flood hazard environment of British Columbia’s Lower Mainland from several directions at once: rising relative sea level, intensifying atmospheric rivers, amplified storm surge, and shifting precipitation regimes. The November 2021 atmospheric river, the costliest natural disaster in the province’s history at an estimated

\$10.6–17.1 billion in total economic losses (Lee & Parfitt, 2022) and attributed as at least 60 percent more likely under anthropogenic warming (Gillett et al., 2022), demonstrated that compound events can overwhelm multiple infrastructure systems simultaneously, and atmospheric-river landfalls in the Fraser basin are projected to increase nearly fourfold by the end of the century (Curry et al., 2019). The region’s highest-value critical infrastructure stands on a subsiding Holocene delta behind a dike system built toward a design standard, which creates the central epistemic difficulty this paper addresses: where protection works, the outcome record is silent about the catastrophic tail, so the risk that matters most is precisely the risk that observation alone cannot estimate.

Research quantifying this risk for Canadian infrastructure remains scarce. A recent systematic review found no probabilistic risk assessments for Canadian infrastructure that consider how multiple hazards will change under future climate scenarios (Sharmin et al., 2026). The Canadian critical-infrastructure flood literature that does exist characterizes resilience rather than probabilistic monetary loss: for Metro Vancouver, infrastructure vulnerability and interdependency have been characterized through structured expert judgments for a seismic scenario (Chang et al., 2014); for Toronto, critical-infrastructure resilience to flooding has been quantified through a disruption response-curve approach over a set of deterministic events up to the 0.3% annual exceedance probability (Murdock et al., 2018). The closest Canadian stochastic flood-loss precedent is an object-based Monte Carlo damage framework applied to building stock (Bryant & Davies, 2022), and the only Canadian application of the CLIMADA platform to date assessed flood risk for generic buildings and utility lines in Halifax, without compound flooding and without aggregate loss exceedance curves across multiple infrastructure sectors (Ryan & Bristow, 2026). For the Fraser estuary itself, the compound flood hazard has recently been characterized through a seasonal statistical–hydrodynamic framework of scenario ensembles (De Jong, 2026); that assessment resolves the hazard and its seasonal structure, not monetary loss, and carries no exposure, vulnerability, or loss exceedance layer. At the global scale, hazard products now provide return-period flood layers for any location, including a 30 m global model spanning fluvial, pluvial, and coastal perils under climate scenarios (Wing et al., 2024); these products model the perils as separate hazard layers and do not carry a joint dependence structure among the flood drivers, which is the structure that governs the compound tail of an estuary. The gap this paper fills is therefore an intersection: an assessment that is simultaneously probabilistic, event-based, compound coastal–fluvial, and resolved to individually identified critical infrastructure assets with aggregate loss exceedance curves.

This paper delivers that assessment for the Fraser River delta. Because the Fraser mainstem is effectively unregulated while its floodplain is heavily diked, the assessment is built in the driver domain: a 600,000-event stochastic catalog samples discharge, non-tidal residual, and physically reconstructed tidal forcing under bi-seasonal copulas selected by formal goodness-of-fit testing, across three anchored sea-level tiers; a reduced-physics hydrodynamic model, evaluated at seven federal gauges, converts drivers into water surfaces with the dike system represented explicitly as HRDEM-derived crest geometry; a Gaussian-process surrogate extends the hydrodynamics to catalog scale; and a custom non-residential damage-function library prices the resulting depths over a census of 376 assets on the CLIMADA platform, yielding year-loss tables, loss exceedance curves, and standard risk metrics per tier, complemented by a named 1894-class design scenario priced outside the probabilistic

framework.

The paper makes three contributions. First, the assessment itself: probabilistic compound flood loss metrics at asset resolution for the critical infrastructure of a major Canadian delta, with the frequency conventions, valuation envelope, and verification lineage documented end to end. Second, a two-component decomposition of compound tail risk that separates the contribution of statistical dependence between the drivers from the contribution of non-linear hydrodynamic coupling within the estuary, extending the attribution framework of Bevacqua et al. (2019); both experiments are executed in full, bounding the role of statistical dependence and revealing a seasonally asymmetric hydrodynamic interaction. Third, a transparency practice: the chain’s five systematic bias mechanisms are declared with their signs and quantitative bounds rather than corrected or softened, so that the published metrics constitute an explicit lower bound with respect to those five mechanisms, with the two-sided uncertainties (footprint sampling, valuation) bracketed rather than absorbed.

Companion analyses that build on the sealed outputs of this assessment (extraction verification, financial-risk translation, and damage-function epistemic uncertainty) are deliberately not developed here. The remainder of the paper is organised as follows. Section 2 establishes the study area and its hydroclimatic regimes. Section 3 documents the data and methods, from the event catalog through the risk conventions to the decomposition design. Section 4 reports the portfolio metrics, concentration structure, and named scenario, including the results of the compound decomposition. Section 5 states the declared systematic bias and limitations, and Sections 6 and 7 discuss and conclude.

2. Study area and hydroclimatic setting

The study area comprises Vancouver’s coastal corridor and the Fraser River delta and fore-shore, from the downtown waterfront and False Creek southward to Roberts Bank and Delta-port, and eastward along the diked lower mainstem to Mission. This corridor concentrates the region’s highest-value critical infrastructure in a narrow band exposed simultaneously to coastal storm surge, riverine flooding, and rising relative sea level. The geomorphological setting is central to the hazard: the Fraser River delta is a Holocene depositional landform undergoing continuous compaction, with natural subsidence of 1–2 mm/yr and anthropogenic subsidence of 3–8 mm/yr decaying with a time constant of approximately 20 years (Mazzotti et al., 2009). The delta plain sits at or near mean sea level, with some areas already up to 1 m below (Dashtgard et al., 2024), and ongoing subsidence averaging -0.4 mm/yr has been confirmed across the delta (Ohenhen et al., 2026). Subsidence augments relative sea level rise by 40–130% (Mazzotti et al., 2009), a factor excluded from provincial guidelines that prescribe 0.5 m by 2050 and 1.0 m by 2100 (BC Flood Hazard Guidelines, 2016).

The Fraser River drains approximately 230,000 km² and is the largest river by discharge on the Pacific coast of Canada. Two gauging stations anchor the quantitative basis of this study: Hope (08MF005, operating since 1912), which integrates the runoff of the upper and middle basin, and Mission (08MH024), approximately 85 km upstream of the river mouth, which defines the upstream boundary of the hydrodynamic model. The discharge marginals of the stochastic catalog are fitted to seasonal annual maxima at Mission (1965–2024); the longer Hope record provides historical context, including the 1894 and 1948 floods of record, but does not enter the fitted marginals. Reconstructions place the June 1894 freshet, the

flood of record, near 17,000 m³/s at Hope; the June 1948 freshet, measured at approximately 15,200 m³/s at Hope, breached the dike system at multiple points and remains the design benchmark of the provincial protection system. No event since 1948 has approached these magnitudes, and paleoflood evidence indicates that freshets of 1894-class magnitude have recurred over the late Holocene, so the instrumental record samples a hazard distribution whose upper tail is real rather than hypothetical (Brice et al., 2021).

Two structural properties of the system shape the methodology. First, the Fraser mainstem was never dammed, so the discharge series is, to a close approximation, a naturalized record of the hydroclimatic driver rather than of a managed outcome; the sole qualifier is the Nechako diversion (since 1952), which depresses mainstem discharge at Hope by approximately 4 to 6 percent, a bias that is conservative for the present purpose (Mundie & Bell-Irving, 1986). Second, the Lower Mainland reach is heavily diked, and dikes alter outcomes without altering drivers: drivers can therefore be sampled statistically from records that protection works cannot filter, while the protection works are represented explicitly inside the hydrodynamic model.

The extreme events of the basin decompose into three mechanistically distinct regimes. Regime A, the snowmelt-dominated mainstem freshet peaking in May–June, produced every major mainstem event of the record and is the primary in-scope mechanism, reaching 8,000 to beyond 13,000 m³/s at Mission. Regime B, atmospheric-river-driven tributary flooding between October and December, produced the November 2021 event, the costliest natural disaster in British Columbia’s history at an estimated \$10.6–17.1 billion (Lee & Parfitt, 2022), attributed as at least 60 percent more likely under anthropogenic warming (Gillett et al., 2022); it enters the present scope partially, through its coastal expression in the storm-season dependence structure. Regime C, emerging hybrid rain-on-snow events, is projected to grow as the rain–snow partition shifts upslope (Musselman et al., 2018), with approximately 45 percent of the basin transitioning to rain-dominated runoff by the 2080s under a high-emission scenario (Islam et al., 2019); it is out of scope for the present catalog and is carried as a documented limitation.

Fitting extreme-value distributions to observed flood outcomes carries a structural bias in heavily protected systems, termed here survivorship bias: behind functioning dikes, extreme driver realizations leave no inundation outcome, so the outcome record systematically underrepresents the hazard the drivers imply, most severely in the upper tail. The driver-domain construction adopted in Section 3 avoids this conditioning: the statistical model is fitted to the drivers themselves (discharge at Mission and non-tidal residual at Point Atkinson), and the protection works are represented explicitly as weir structures with HRDEM-derived crest elevations (1 m sampling along dike centrelines), so the transition from contained to catastrophic outcomes is computed by the physics for every sampled driver combination.

The catastrophic envelope of the system is compound: the joint occurrence of elevated Fraser discharge, elevated coastal water level, high tide, and rising relative mean sea level, in the multivariate compound-event typology of Zscheischler et al. (2018, 2020), with sea level rise acting as a systematic amplifier of compound flood potential in estuaries (Wahl et al., 2015; Moftakhari et al., 2017). The historical record contains no catastrophic coincidence of a major freshet with a major storm surge, but reading this absence as physical impossibility would repeat the survivorship error: the measured dependence between the drivers is weak but non-zero in both seasons (Kendall’s $\tau = 0.130$ in the freshet season and

$\tau = 0.084$ in the storm season), consistent with the comparatively weak surge–discharge dependence documented for Pacific coasts (Couasnon et al., 2020; Jalili Pirani & Najafi, 2023), and independently corroborated for this estuary by De Jong (2026), whose peaks-over-threshold analysis finds the dependence between discharge and the non-tidal residual statistically indistinguishable from zero at extreme thresholds in all three seasons, with weak dependence emerging only at moderate thresholds. The scope of the assessment is therefore the probabilistic, compound, mainstem-and-tide catastrophic flood risk of the dike-protected lower mainstem and delta west of Mission, evaluated on an inventory of 376 critical infrastructure assets enumerated exhaustively within five infrastructure categories defined ex ante (Section 3.2). Figure 1 presents the study domain, the model boundaries, and the asset inventory.




Figure 1: Study domain and model setup: the diked lower Fraser mainstem and delta west of Mission, the hydrodynamic model domain with its water-level and outflow boundaries, and the 376 critical infrastructure assets of the census.

3. Data and methods

3.1. Chain overview

The architecture follows from the epistemology of Section 2: because the drivers are observable and effectively naturalized while the outcomes are protection-filtered, the hazard

is built in the driver domain and converted to outcomes by physics. The chain has five links (Figure 2). First, a bi-seasonal copula catalog of 600,000 synthetic events is generated over discharge, non-tidal residual (NTR), tidal sampling instant, and sea level rise (SLR) tier (Section 3.3). Second, a multi-gauge-evaluated reduced-physics hydrodynamic model converts drivers into maximum water surfaces with the dike system represented explicitly (Section 3.4). Third, a hybrid surrogate extends the hydrodynamics from the simulated design events to the full catalog (Section 3.5). Fourth, an exposure module fixes what is at risk and at what replacement value (Section 3.2), and a vulnerability module fixes how depth becomes damage through custom impact functions (Section 3.6). Fifth, a risk engine multiplies hazard, exposure, and vulnerability into per-event losses, year-loss tables, and loss exceedance curves under documented frequency conventions (Section 3.7).




Figure 2: The five-link production chain. Parallel modules are separated by ||; sequential dependencies by →.

The chain delivers three hazard products on one common vertical rule: a present-day baseline resting on 122 weighted direct simulation events, three sea-level-rise tiers of 200,000 seasonal events each, yielding 100,000 annual maxima per tier, and a named class-1894 design scenario that carries no frequency coordinate by construction. The design of the compound decomposition, which attributes tail risk separately to statistical dependence and to hydrodynamic coupling, closes the methods (Section 3.8). Throughout, the state described is the sealed production state of the chain; design elements whose production runs do not yet exist are identified as such where they appear.

3.2. Asset inventory and exposure valuation

The exposure universe comprises 376 lifeline facilities enumerated exhaustively within five infrastructure categories defined ex ante; reported average annual loss (AAL) and probable maximum loss (PML) are a lower bound with respect to airport and other excluded infrastructure sectors. The five categories are: energy infrastructure (23 BC Hydro substations), maritime and port operations (31 assets), mass transit (82 assets), civic and emergency facilities (206 assets), and water and wastewater infrastructure (34 assets), of which 111 assets lie on Fraser delta sediments and 265 in the broader Lower Mainland. The inventory was assembled through a hybrid strategy combining authoritative top-down identification from agency records with bottom-up discovery from OpenStreetMap via the Overpass API, followed by a structured four-version pipeline of cross-validation, false-match correction, operator filtering, and QGIS-based quality control against satellite imagery.

Assets are classified as compact or extended by the relationship between footprint and hazard grid resolution. Compact assets enter as single georeferenced points. Extended assets

(port terminals, large transit and maintenance centres, the Iona Island Wastewater Treatment Plant, major works yards) are disaggregated into a regular grid of points at a spacing no coarser than the hazard resolution, following the points-on-polygon precedent of Ryan & Bristow (2026) and the CLIMADA `lines_polys_handler` module of Mühlhofer et al. (2024c); the parent value is allocated across the points and losses are re-aggregated to the parent asset after impact computation. One documented exception is retained: the Annacis Island Wastewater Treatment Plant enters as a single facility-level exposure referenced to the regulated-pond extraction cell of the production chain, the physically distinguished point of the site because its permanent, hydraulically controlled water level anchors the vertical rule and the permanent-water floor; because the 36 ha footprint spans substantial within-site hazard variation, Section 4.2 interrogates this exception explicitly through a full-footprint sensitivity bracket rather than leaving it implicit.

Each asset is valued parametrically at replacement cost, as construction area times a unit cost by typology, with valuation resolved at the level of 32 areal subtypologies. Unit costs are compiled exclusively from publicly available sources, including agency capital plans, regulatory filings, reference values by voltage level for substations, and peer-reviewed expert-derived cost studies (PNNL, 2016; BC Hydro & British Columbia Utilities Commission, 2025; Williams et al., 2025; Martello et al., 2023; Altus Group, 2025); proprietary commercial sources are excluded to preserve reproducibility. All unit costs are documented as ranges (lower, central, upper) rather than point values. The full unit-cost matrix and its provenance are provided in the Supplementary Information.

The exposure census comprises 376 assets with a central replacement value of CAD 55.61 billion. Five business-interruption-only rows (Emergency Operations Centre contents, CAD 953 million) carry no damage curve and are excluded from the direct-damage exposure, which totals CAD 54.66 billion; companion analyses refer to this direct-damage exposure as the insured-value proxy.

Business interruption is outside the scope of the results reported here. The assessment documents a two-component multiplier design (direct service disruption and cascading disruption of dependent services), with category-level ranges anchored in the empirical literature and a network-based cascade framework specified as its validation pathway; the design and its ranges are reproduced in the Supplementary Information as documented inputs consumed by companion analyses, not as results of this paper. The five business-interruption-only exposure rows (CAD 953 million) enter the risk chain only through that companion pathway.

3.3. Stochastic compound event catalogue

Generalised extreme value (GEV) marginals are fitted to seasonal annual block maxima of the two drivers: Fraser discharge at Mission (53 freshet-season and 33 storm-season maxima, 1965–2024) and the non-tidal residual at Point Atkinson (70 years of water level data). The fitted shape parameters are $\xi = -0.154$ (freshet discharge, a bounded Weibull-type tail with finite upper support at 19,944 m³/s), $\xi = -0.013$ (freshet NTR), $\xi = +0.090$ (storm discharge, the only unbounded tail), and $\xi = -0.684$ (storm NTR). The bounded freshet tail is a statement about the fitted marginal, not the physical process: the Mission fitting sample contains neither the 1894 nor the 1948 flood of record, and the 1894-class discharge referenced to Mission ($\sim 18,900$ m³/s; NHC 2008a) carries a nominal analytic return period of order 10⁶ years under this marginal. The consequences of this bound, and of the sampling

ceiling introduced below, are treated by the named design scenario of Section 4.3 and declared in Section 5.

Dependence between the drivers is modelled with bi-seasonal copulas, calibrated over an extended candidate set of ten families (Archimedean, survival rotations, elliptical, and Tawn extreme-value), motivated by alignment with the most recent benchmark for Canadian coasts (Jalili Pirani & Najafi, 2023). Goodness-of-fit uses the Cramér–von Mises S_n statistic of Genest et al. (2009) with 500-replicate parametric bootstrap, and the pre-specified selection rule, applied identically to both seasons, is minimum AIC among copulas passing goodness-of-fit at $p > 0.05$. In the freshet season the Gaussian copula is selected: it is the unique family passing goodness-of-fit ($p = 0.236$) and additionally has the lowest AIC, structurally consistent with the selection at Tofino, the closest Pacific-coast benchmark site. In the storm season the Clayton copula is retained ($p = 0.248$), whose weak lower-tail dependence is physically coherent with the synoptic structure of British Columbia winters. The full selection diagnostics (AIC separations, tail-dependence coefficients, and bootstrap confidence intervals for all ten families) are provided in the Supplementary Information. The selected families carry no upper-tail asymptotic dependence, and the selection operates on the body of the joint distribution, with limited power against tail alternatives at these sample sizes; the dependence contribution estimated in Section 3.8 is accordingly conditional on this selection. The conditionality is empirically supported rather than merely declared: the independent peaks-over-threshold analysis of De Jong (2026) for the same estuary finds the discharge–surge dependence vanishing, not strengthening, as event thresholds rise.

Tidal forcing is sampled physically rather than statistically: for each event a sampling instant t_0 is drawn uniformly within the 70-year UTide harmonic calibration window and the harmonic prediction is reconstructed over a 14-day (336 h) event window centred on t_0 . The 14-day window guarantees at least one full spring–neap cycle in every event (the M_2 fortnight has a period of 14.77 days; Pugh & Woodworth 2014) and is conservative with respect to the 3–5-day joint-event windows adopted on multi-driver coasts (Eilander et al., 2023a; Maduwantha et al., 2026). Drawing the causal instant rather than a tidal elevation preserves the spring–neap, perigean, and seasonal modulations of the natural signal and produces a tidal predictor statistically independent of (Q, NTR) by construction, a property required both by the surrogate of Section 3.5 and by the independence-copula counterfactual of Section 3.8.

Each sampled event is replicated across three sea level rise tiers by varying only the additive SLR offset, producing 600,000 events in total ($100,000 \text{ years} \times 2 \text{ seasons} \times 3 \text{ SLR tiers}$). Generator uniforms are drawn on the open interval $(0.001, 0.999)$, imposing a rank-sampling ceiling: the catalog maximum freshet discharge ($16,086 \text{ m}^3/\text{s}$) is the empirical realisation of the ceiling quantile $F_Q^{-1}(0.999) = 16,089.6 \text{ m}^3/\text{s}$, an analytic ~ 990 -year event of the fitted marginal, and no draw can exceed it. Catalog convergence is verified empirically through a two-stage test ($10,000$ versus $100,000$ years, and $100,000$ versus $500,000$ years), with all tail metrics agreeing within 0.39% on total water level, the primary risk driver, and within 0.011% on the net tidal residual P99.

Sea level rise enters as a per-tier additive offset to the coastal boundary forcing. Table 1 maps the tier labels used throughout this paper and its companion analyses to their operative offsets and source anchors. The tiers are relative sea-level planning levels with citable provenance, not trajectory forecasts: the 0.77 m value is simultaneously the IPCC AR6 median for

SSP5-8.5 and approximately the upper bound of the likely range for SSP2-4.5 (0.44–0.76 m; Fox-Kemper et al. 2021), so the tier structure is reachable across scenarios, and the delta subsidence component of Tier C is independent of any emission scenario (Mazzotti et al., 2009; Dashtgard et al., 2024).

Table 1: Sea-level tier mapping: label, operative offset, and source anchor. Labels SLR_A, SLR_B, and SLR_C are used identically by the companion analyses.

Label	SLR level (2100 planning horizon)	Source anchor	Role
SLR_A	+1.00 m	BC Flood Hazard Guidelines (BC Flood Hazard Guidelines, 2016)	Headline tier
SLR_B	+0.77 m	IPCC AR6 median (Fox-Kemper et al., 2021)	Global benchmark
SLR_C	+1.50 m	Delta-adjusted: AR6 + subsidence (Mazzotti et al., 2009; Dashtgard et al., 2024)	Delta-adjusted tier

Because the generator’s rank labels and the fitted marginals answer different questions, all frequency-based results carry a dual-clock convention. The event carrying rank label T_{rank} corresponds to the nominal uniform mapped affinely onto (0.001, 0.999), giving the analytic driver-clock conversion

$$T_{\text{analytic}} = \frac{1}{0.001 + 0.998/T_{\text{rank}}}, \quad T_{\text{analytic}} \xrightarrow{T_{\text{rank}} \rightarrow \infty} 1,000. \quad (1)$$

The two clocks agree within 10% up to $T_{\text{rank}} \approx 100$ and the analytic clock saturates at 1,000 years: no rank label, however large, corresponds to an analytic rarity above 1,000 years. Beyond the saturation, rank-clock positions remain meaningful as internal orderings of the simulated world and are flagged as such wherever they appear.

3.4. Hydrodynamic modelling and validation

Drivers are converted into water surfaces by SFINCS (Super-Fast INundation of CoastS) v2.3.2, a reduced-physics hydrodynamic solver that computes compound flooding from fluvial, tidal, and coastal processes at a computational cost approximately three orders of magnitude lower than full-physics models (Leijnse et al., 2021), solving simplified momentum and continuity equations with an explicit first-order scheme (Bates et al., 2010). The model domain covers the delta from the Strait of Georgia to upstream of Mission on a 50 m grid of $1,499 \times 900$ cells (EPSG:26910), with a 5×5 subgrid per cell yielding an effective 10 m resolution, and was constructed with HydroMT-SFINCS v1.1.0 (Eilander et al., 2023b).

A single elevation source is used: the Ocean Networks Canada Canadian Continental Shelf Program (ONC CSSP) 2019 topo-bathymetric digital elevation model, 15 tiles at 3 m resolution in the CGVD2013 vertical datum, which provides real bathymetric depths for the river channel (thalweg elevations from -12 m at Steveston to -27 m near Mission) rather than the water surface that LiDAR-based models register. The dike system is represented explicitly as 218 weir segments with HRDEM-derived crest elevations (1 m sampling along dike centrelines with a buffered fallback), confirmed by direct count of the production model structure file; a sensitivity analysis showed the weir configuration does not affect in-channel validation metrics, the weirs mattering for lateral floodplain inundation.

The validation reported here follows a methodological audit (May 2026) that invalidated an earlier single-gauge validation on two independent grounds (an incorrect vertical datum conversion of the observed series, and model extraction at incorrect gauge coordinates); the superseded metrics are not cited. Remediation proceeded through three minimal-difference

iterations (A-Min, B-Min, C-Min), each isolating an attributable change, with a pre-specified fallback to a re-scoped fitness-for-purpose statement if tolerance bands were not met. The production configuration C-Min improved on A-Min at all seven gauges with no degradations.

C-Min is validated against the June 2012 freshet at seven federal gauges along the lower mainstem. It carries a characterized, systematic, conservative misfit: mean absolute peak bias of 0.76 m and multi-gauge root-mean-square error (RMSE) of 0.87 m, with six of seven gauges under-predicting the observed peak. Whonock under-predicts by 0.128 m and passes the chainage-dependent tolerance band (± 0.260 m); four non-excluded gauges under-predict by 0.74 to 1.01 m; Pitt River under-predicts by 1.50 m but is excluded from the classification verdict as a tributary station; and Mission over-predicts by 0.329 m. Under the global Defra/EA estuarine classification compiled by Williams & Esteves (2017), the multi-gauge RMSE exceeds the 0.30 m threshold and the configuration is classified as Unsatisfactory for design-grade water-level prediction. This classification is stated explicitly rather than softened.

The configuration is adopted under an explicit fitness-for-purpose statement: the misfit is directional (under-prediction growing monotonically with water level at the tidal-dominated gauges), it is measured per gauge, and it is treated at the loss-aggregation stage rather than by hydraulic tuning, preserving the physical consistency of the frozen model. Because under-predicted water levels propagate to under-predicted inundation depths and losses, the uncorrected results of Section 4 constitute a defensible lower bound on losses with respect to the characterized hydrodynamic bias. The full bias treatment, its five constituent mechanisms, and the declared limitations are the subject of Section 5. Figure 3 summarises the validation verdict and the declared limitations of the production configuration.

A second, independent reference complements the June 2012 event validation: the named class-1894 scenario of Section 4.3 includes an anchor check against the official provincial design profile, in which the model brackets the design water levels at Steveston (+0.90 m), New Westminster (−1.00 m), Port Mann (−1.61 m), and Mission (+0.49 m). These anchors are distinct from the June 2012 per-gauge biases quoted above: they compare the model against the design profile of the protection system under an extreme synthetic forcing, not against an observed event, and their interpretation is developed with the scenario itself.

3.5. Surrogate extension

Running the hydrodynamic model for all 600,000 catalog events is computationally infeasible, so the chain is extended by a hybrid surrogate following the Low-fidelity, Spatial analysis, and Gaussian Process learning (LSG) framework of Fraehr et al. (2022, 2023). A design set of exactly 300 events (160 freshet, 140 storm) is selected by stratified sampling over five joint-return-period bins crossed with the six (season, SLR tier) strata, with importance weights summing to the 600,000 catalog events, and each design event is simulated directly with the production hydrodynamic configuration.

The surrogate architecture was selected by a factorial sweep comparing an emulator (Gaussian Processes on empirical orthogonal function coefficients of the high-fidelity field) against a corrector (learning the high-minus-low-fidelity residual). The emulator outperforms the corrector at every point of the design grid (marginal-mean depth RMSE 0.071 m against 0.128 m), and the selected configuration uses the predictor set $\{wl_{\max}, Q, SLR\}$, where $wl_{\max}$ is the event-peak boundary water level, with 99% variance truncation ($k \approx 25$ modes) and




Figure 3: Production hydrodynamic configuration: per-gauge validation verdict against the June 2012 event and declared limitations, including the dike representation as weir structures and the fitness-for-purpose adoption under the Defra/EA classification.

Matérn-5/2 kernels with automatic relevance determination. Removing discharge from the predictor set degrades accuracy markedly for both architectures, so discharge is informative for the spatial structure of the field even where the boundary water level carries the dominant signal.

Surrogate accuracy is evaluated by five-fold cross-validation over the 300 design events, stratified by season and SLR tier, achieving a cross-validated per-asset on-land inundation-depth RMSE of 0.053 ± 0.004 m, restricted to the 361 on-land assets (15 over-water structures are treated separately). This emulation error is roughly an order of magnitude below the characterised bias of the high-fidelity model it learns from; its contribution in depth space is therefore secondary, although threshold effects can amplify local depth errors in loss space (Section 4.4); Fraehr et al. (2024) showed that LSG outperforms purely data-driven surrogates precisely in the tail-extrapolation regime relevant here. The production output is a mean-depth and a standard-deviation matrix, each $600,000 \times 376$, with a mean propagated uncertainty of 0.020 m per asset depth.

The present-day baseline is not a surrogate product: it rests on a dedicated stratified campaign of 122 directly simulated events carrying stratum sampling weights. In total the chain rests on 440 direct high-fidelity simulations: 300 stratified design events, 122 baseline campaign events, 3 class-1894 scenario variants, and 15 targeted verification events. Experiment 2 adds 60 counterfactual SFINCS simulations used exclusively for the compound-interaction analysis, outside the sealed production chain. The hazard quantity consumed by the risk module is in all cases the excess-over-floor depth matrix, in which the permanent water floor is subtracted so that only inundation in excess of it produces damage.

3.6. Vulnerability: BC-adapted damage functions

Depth–damage functions for the infrastructure typologies of this study cannot be adopted directly from existing sources: standardised databases such as HAZUS-MH are calibrated to United States building stock and ship their electric-power, water, and pipeline facility curves in a disabled mapping-only state that produces no loss by default (HAZUS-MH v7.0, 2025); continental functions lack infrastructure sub-type differentiation (Huizinga et al., 2017); and infrastructure-specific functions are scarce, geographically distant, or demonstrated to understate damage (Williams et al., 2025). The claims-based finding of Wing et al. (2020), that conventional depth–damage curves exhibit severe deficiencies, is carried as a directional caution on any single deterministic curve.

The library is therefore built by a two-step process: (i) compilation from open sources into a source-of-record base, and (ii) literature-guided adaptation to the British Columbia context with documented rationale, following the template of Williams et al. (2025). Three principles govern the specification. Every function is an explicit closed-form equation or complete numeric array. Structural and contents values are separated only where the source supplies distinct curves, with losses re-aggregated to the parent asset. Every function carries an evidence tier (a literature depth–damage array; a fragility or binary threshold; or an expert-judgement proxy with the gap named), so the weakest-evidence cases are labelled rather than hidden. In the CLIMADA representation each function returns a mean damage ratio (MDR) as the product of two interpolated vectors,

$$\text{MDR}(I) = \text{PAA}(I) \cdot \text{MDD}(I), \quad (2)$$

where I is flood depth, MDD the mean damage degree, and PAA the fraction of assets affected, making the two factors individually meaningful and directly perturbable in sensitivity analysis.

The principal quantitative bases of the compilation are the HAZUS tables of Pollack (2023), the continental functions of Huizinga et al. (2017), the expert-elicited curves of Williams et al. (2025), and the rail-transit curves of Martello et al. (2023), with supporting sources including the global infrastructure database of Nirandjan et al. (2024). The adaptation step documents, for every adapted function, the delta between the source curve and the production curve and its rationale.

The production library comprises 25 impact functions covering the non-residential typologies of the inventory, assigned through a documented subtypology-to-impact-function crosswalk that keeps the valuation granularity (32 subtypologies) and the damage-function granularity reconcilable; the source provenance of every curve is sealed under a three-part key (occupancy code, source table, description), since HAZUS damage-function identifiers are not unique across source tables, and 109 provenance keys are recorded in the compilation report. The complete function library, the crosswalk, and the adaptation deltas are provided in the Supplementary Information.

Application of the functions to the inventory required a set of data-conditioning decisions executed at the exposure build and recorded explicitly: normalisation of treatment-plant asset-type strings; inclusion and deduplication hygiene in the water and wastewater category (removal of physically inadmissible discoveries, reconciliation of fragmented polygons and one cross-category duplicate); assignment of open-process versus closed-process treatment curves by physical enclosure; a single life-stage snapshot per facility where an existing plant and its replacement under construction coexist; and routing of loss-of-function consequences (Emergency Operations Centres, treatment-plant public-health consequences) to the business-interruption pathway rather than to a damage curve.

3.7. Risk computation and conventions

The impact engine receives the excess-over-floor depth matrices, the 494-row direct-damage exposure (CAD 54.66 billion central), and the 25 production impact functions, and is implemented on the CLIMADA platform (Aznar-Siguan & Bresch, 2019). The persisted primitive is not the loss but the mean damage ratio per event and exposure row, stored as a sparse matrix; losses derive from it by multiplication with the row values, so the three valuation variants (low, central, upper) share one hazard-times-vulnerability computation and the valuation envelope remains exactly consistent across every downstream metric. Impact functions are evaluated with $MDR(0) = 0$ pinned and guarded at the engine boundary, with any recurrence of a nonzero damage ratio at zero depth a fatal condition of the engine's selftest; the engineering detail of this guard and of the chunked processing of the tier sets is documented in the Supplementary Information.

Every production run carries a verification pass in which the CLIMADA computation is reproduced by an independent NumPy implementation sharing only the input artefacts; the maximum relative divergence observed in production is of the order of machine epsilon (2.05×10^{-16} for the baseline set), verifying that the engine applies exactly the documented curves to exactly the documented depths.

Frequency conventions are documented per product. For the tier catalogs, the year identity of an event is recovered arithmetically and a structural guard verifies that the 600,000 events resolve into six blocks of exactly 100,000 unique years before any frequency is computed; the annual loss of a year is the maximum of its two seasonal event losses (the occurrence-exceedance convention standard in catastrophe modelling), the 100,000 annual maxima per tier are ranked into the loss exceedance curve, and the per-tier year-loss tables are persisted products of the chain. Tier AAL is accordingly the mean of the 100,000 annual maxima (occurrence basis); the aggregate-basis AAL under the seasonal-sum convention is reported in the dual-convention table of the Supplementary Information (aggregate-to-occurrence ratios of 1.20, 1.23, and 1.66 at the three tiers, and 1.004 at baseline). Because the annual maximum is bounded above by the annual sum, the occurrence-basis tier AALs, and every cross-tier multiplier built on them, are conservative relative to the aggregate convention, one further element of the declared lower-bound orientation. For the present-day baseline, frequencies are recovered by stratum weighting of the 122 campaign events (weights summing to the 200,000 events of the stratification frame, with twelve zero-weight guard events), and the baseline AAL is computed under the event-sum convention; because the campaign was stratified toward the tail, frequent-event coverage is sparse by design and the baseline AAL carries that sampling caveat. All frequency-based metrics are reported under the dual-clock convention of Equation (1). Loss exceedance curves are constructed at three aggregation levels (individual assets, infrastructure categories, full portfolio), with the aggregate curve summing losses within each event before ranking, so spatial correlation between assets is implicit and event-driven.

3.8. Design of the compound decomposition

The contribution of compound behaviour to tail risk is decomposed into two components, following the attribution framework of Bevacqua et al. (2019). Compound flooding studies typically report a lumped compound effect; Bevacqua et al. (2019) introduced a decomposition of dependence against marginals, and the present design extends it to a decomposition of statistical dependence against hydrodynamic coupling, an extension enabled by the use of a process-based hydrodynamic model rather than a purely statistical framework. The two experiments below are specified and executed in full; Section 4.4 reports the results.

Experiment 1: statistical-dependence contribution (independence counterfactual by permutation). The independence counterfactual follows the dependence-versus-independence contrast of Zscheischler and Seneviratne (2017), implemented as a block permutation that preserves the empirical margins and the year structure of the catalogue exactly: within each season, the surge-and-tide block (NTR and the tidal sampling instant t_0) is re-paired at random against the discharge draws at the level of the 200,000 base events, and the same re-pairing is propagated to the three tiers of each base event. The permutation imposes the product copula on the exact empirical marginals (Kendall’s τ collapses to within ± 0.005 of zero in both seasons) while leaving every marginal draw, every year-season block, and every tidal instant untouched; three permutation seeds are run to expose the Monte Carlo variability of the counterfactual itself. The counterfactual catalog is evaluated with the same trained LSG surrogate of Section 3.5, with no retraining: because the permutation only re-pairs existing draws, each event’s predictors ($w_{\max}$, Q , SLR offset) coincide with sealed production values, the multiset of $w_{\max}$ is identical to production by construction, and the out-of-support

fraction equals the production value exactly. The difference between the production loss exceedance curve and the independence-copula curve quantifies the statistical-dependence contribution:

$$\Delta L_{\text{stat}}(p) = L_{\text{compound}}(p) - L_{\text{indep}}(p) \quad (3)$$

where $L(p)$ is the loss at exceedance probability p .

Experiment 2: hydrodynamic-interaction contribution (four-leg single-driver design). A four-leg differences-in-differences design isolates the non-linear interaction that the hydrodynamic model resolves between extreme discharge and extreme NTR within the estuary, run entirely through the full SFINCS model rather than the surrogate. Selection is restricted to the 300-event high-fidelity design pool, whose extreme corner reaches the 99.7th to 99.997th percentile of the corresponding catalog driver, so that the 20 compound legs reuse sealed production runs: the 10 most extreme events in Q (stratified 4/3/3 across tiers) and the 10 most extreme in NTR. Each event is complemented by three counterfactual forcings that preserve its tidal instant, tier, and sea-level offset exactly: a Q -only leg (NTR at its seasonal median), an NTR-only leg (Q at its seasonal median), and a both-at-medians leg that measures the background loss the tier produces without either driver extreme. The net hydrodynamic interaction for each tail event i is

$$I^{(i)} = L_{\text{compound}}^{(i)} - L_{Q\text{-only}}^{(i)} - L_{\text{NTR-only}}^{(i)} + L_{\text{med-med}}^{(i)} \quad (4)$$

with the fourth leg preventing the double counting of the background term that a three-leg subtraction would incur; the identity $L_{\text{compound}} = B + \Delta_Q + \Delta_{\text{NTR}} + I$ holds exactly by construction. A positive I indicates that the joint hydrodynamic action exceeds the linear superposition of the single-driver excesses over background, that is, non-linear amplification. A parallel decomposition applied to station water levels rather than losses separates hydraulic superposition from the non-linearity of the damage chain.

Reporting convention. The two contributions are reported separately rather than as a single compound effect. For tail events (return period above 100 years), the fraction of aggregate losses attributable to statistical dependence and to hydrodynamic coupling is tracked across SLR tiers. The full operational protocol for both experiments, including independence-copula catalog construction with sanity checks, single-driver scenario design, surrogate application procedures, and bootstrap confidence interval computation, is provided in the Supplementary Information.

4. Results

4.1. Portfolio risk metrics and loss exceedance curves by tier

Five standard risk metrics are extracted at each aggregation level: the average annual loss (AAL) and the probable maximum losses (PML) at the 50-, 100-, 200- and 500-year return periods, each quoted under both clocks of Equation (1) and at the three valuation variants, with central values in the text and the low/upper envelope carried by the figures. Table 2 reports the headline metrics; Table 3 reports portfolio losses across the frequency spectrum; Figures 4 and 5 present the loss exceedance curves. The cell-level extraction verification protocol behind these production values, including the override registry through which

six contaminated extractions were identified and corrected, is documented in a companion manuscript; the metrics reported here are the post-verification production state.




Figure 4: Portfolio loss exceedance curves, 376 critical assets, under the three SLR tiers and the present-day baseline. Lower axis: native rank clock; upper axis: analytic driver clock, $T_{\text{analytic}} = 1/(0.001 + 0.998/T_{\text{rank}})$, validity $T \lesssim 500\text{--}1,000$ yr. Shaded bands: valuation envelope (low/upper). Horizontal dashed lines: class-1894 named scenario losses at the matching SLR levels (Section 4.3).

The AAL-to-PML-100 ratios carry the central qualitative reading. Diversified catastrophe portfolios typically show ratios of 0.05–0.3, a practitioner heuristic from catastrophe-modelling practice rather than a published bound; the present-day baseline sits below that band, at 0.018, and this is a finding rather than a defect. The band describes portfolios without uniform protection; a delta whose critical infrastructure stands behind dikes built toward a design standard produces exactly the observed regime: most years contribute essentially nothing (the present-day AAL is CAD 0.374 M against CAD 55.6 B of exposure), and the annual average is carried by rare overtopping events (the largest baseline campaign event produces CAD 98 M). The baseline ratio is the numerical signature of functioning protection. The ratio then rises monotonically with sea level, 0.088 at +0.77 m, 0.189 at +1.00 m, and 0.702 at +1.50 m, tracing the transition from an episodic, tail-dominated





Figure 5: Portfolio loss versus return period, central values with the valuation envelope shaded. Top: native rank clock, full simulated range. Bottom: analytic driver clock, with the saturation at $T = 1,000$ marked (validity $T \lesssim 500-1,000$ yr).

Table 2: Portfolio risk metrics by sea-level tier (central values, CAD M).

	Present day	SLR_B (+0.77 m)	SLR_A (+1.00 m)	SLR_C (+1.50 m)
AAL, occurrence basis	0.374	119.6	350.6	1,726.2
AAL, aggregate basis	0.374	143.6	429.9	2,868.3
PML-100	20.6	1,365.1	1,858.5	2,457.7
PML-500	56.2	1,733.5	1,939.0	2,904.4
AAL/PML-100	0.018	0.088	0.189	0.702

Table 3: Portfolio loss at standard return periods (rank clock, central values, CAD M).

	50 yr	100 yr	200 yr	500 yr
Present day (SLR 0)	9.5	20.6	25.5	56.2
SLR +0.77 m	1,028	1,365	1,403	1,733
SLR +1.00 m	1,800	1,859	1,897	1,939
SLR +1.50 m	2,283	2,458	2,657	2,904

regime to a chronic one in which routine years flood protected assets. The migration survives under either consistent convention pairing (0.018 to 0.69 under the aggregate basis), and the aggregate-to-occurrence AAL ratio of 1.66 at the highest tier is itself a signature of the chronic regime: with both seasonal events producing losses in most years, the annual sum approaches twice the annual maximum. Between the +0.77 m and +1.50 m tiers the AAL multiplies by a factor of approximately 14 (119.6 to 1,726.2 CAD M).

The chronic character of the tier regimes is direct in the year-loss tables: under the occurrence convention of Section 3.7, 99.99%, 99.64%, and 100.00% of the 100,000 simulated years carry positive portfolio loss at the +1.00, +0.77, and +1.50 m tiers respectively (99,991, 99,641, and 100,000 years), against a present-day baseline in which most years produce no loss at all.

The multiplier of approximately 320 between the present-day AAL and the first tier (384 under the uniform aggregate basis, so the occurrence-basis figure is the conservative one) invites scrutiny and has three independent supports. First, dike-protected systems are threshold non-linear: protection is quasi-binary in freeboard, so consuming the margin converts near-null years into loss-bearing years. Second, the named class-1894 scenario of Section 4.3 reproduces the jump independently of the surrogate and of the frequency framework, scaling a single fixed hydraulic event by a factor of approximately 40 under +0.77 m alone. Third, the baseline AAL of 0.374 M carries its documented event-sum convention and the stratified-campaign sampling caveat of Section 3.7, so the two ends of the multiplier are computed under declared, distinct conventions rather than presented as a single homogeneous series.

The frequency axis of these curves carries a declared information-set correction. The official flood-frequency estimates, which incorporate the 1894 and 1948 historical floods through the Expected Moments Algorithm over an extended ~ 162 -year period, sit systematically higher than the fitted instrumental marginal: by approximately 15% at $T = 100$ and approximately 22% at $T = 500$ on the common station and clock. The divergence is attributable to the information set, not to the distribution family. The consequence for the loss exceedance curves is directional and declared: extreme-tail losses are priced at the correct severity but

likely understated in frequency, that is, events of a given magnitude arrive somewhat more often than the fitted curve implies. This is one of the five bias mechanisms of Section 5.

4.2. Concentration and dominant contributors

The tier losses are highly concentrated. Of the 376 assets, 56 carry nondegenerate loss in at least one tier: 22 at +0.77 m, 27 at +1.00 m, and 56 at +1.50 m, a monotonically nesting set consistent with progressive freeboard consumption. The remaining assets are structural zeros of the tier product: high-ground or out-of-corridor sites that the hydrodynamic envelope never wets, verified for the lowland subset by targeted direct simulation campaigns (Section 3.5).

Within the nondegenerate set, two wastewater assets dominate. Annacis Island Wastewater Treatment Plant carries 55.1%, 63.2%, and 83.5% of the tier AAL at +0.77, +1.00, and +1.50 m respectively; Iona Island Wastewater Treatment Plant carries 42.1%, 34.4%, and 11.5% at the same tiers, so the pair jointly carries 97.5% of the headline-tier (+1.00 m) AAL. The concentration structure is not static across tiers: as sea level rises, Annacis progressively displaces Iona in the attribution, because the Annacis threshold analysed below converts that single asset from a tail contributor into a routine one.

Concentration of this magnitude is exactly the class of signal that triggered the extraction verification, so it is not quoted here until interrogated on three grounds. Extraction integrity: the Annacis baseline series shows a regulated permanent water level of 0.237–0.245 m fully neutralised by the permanent-floor rule, with a residual excess of at most 8 mm; the extraction cell is clean. Arithmetic proportionality: the dominance decomposes mechanically as value share times damage-ratio escalation. Annacis enters the exposure as a single facility row valued at CAD 3.89 B, 7.1% of the CAD 54.66 B direct-damage portfolio, and its mean damage ratio over wet events escalates from 0.022 at +0.77 m to 0.067 at +1.00 m and 0.309 at +1.50 m; an independent reconstruction of its AAL share from the sparse damage-ratio matrix reproduces the production attribution. The concentration is the arithmetic of a very large value on a genuinely responding site, not a valuation or curve-assignment error.

The physics: a threshold between +1.00 and +1.50 m. The median tier depth at Annacis is 0.23 m at +0.77 m and 0.25 m at +1.00 m, the signature of the regulated pond neutralised by the floor, with the AAL contribution at those tiers coming from tail events; between +1.00 and +1.50 m the median jumps to 3.12 m. At +1.50 m of relative sea-level rise the island floods in the median annual event: the protection of the site is overtopped routinely rather than exceptionally. The threshold is hydraulically coherent, since delta-front water levels of the +1.50 m tier reach 3.9–4.0 m (GSC) while typical lower-Fraser crest elevations lie at 3.5–4.7 m, and the named scenario confirms the same threshold from an independent product (Section 4.3).

Three consequences are carried forward. First, with one asset carrying more than half of the tier AAL, the aggregate exceedance curve is dominated by the largest individual curve from the shortest return periods. Second, the single-point representation of an extended asset makes Annacis the named subject of the Jensen-inequality sensitivity bracket of Section 5: with 83.5% of the +1.50 m AAL on one polygon extracted at one point, the bracket over its footprint is the sensitivity of the aggregate, and it is executed below. Third, the threshold statement is conditional on the static, overtopping-only representation of flood defences; structural failure modes would move the threshold earlier, never later, preserving the conservative orientation of the chain.

The footprint bracket, executed by re-evaluating the surrogate hazard cell by cell over the 36 ha plant footprint (118 model cells; details and full table in the Supplementary Information), bounds the effect of within-footprint hazard sampling on the aggregate. The bracket is wide: sampling at the wettest riverfront cell, which is inundated in 98.9% of catalog events, would raise the +0.77 m portfolio AAL from 119.6 to 1,609 CAD M, while sampling at the driest admissible cell, wet in 2.2% of events, would lower it to 59 CAD M; the production convention, anchored at the regulated pond with its permanent-water floor, sits between the extremes on physically motivated grounds rather than at either. The concentration structure and the absolute AAL level are therefore sensitive to the sampling convention, and the timing of the Annacis threshold reported above is a property of that convention. The regime migration is not: under the driest admissible sampling the portfolio AAL remains 160 to 1,200 times the present-day value across the tiers, the AAL-to-PML-100 ratio remains far above the present-day 0.018, and the deep tail of the +1.50 m tier is entirely insensitive to the bracket, with PML-100 and PML-500 unchanged across all three variants. The head-line regime shift survives the sampling question; the attribution detail does not, and both statements are carried forward as written.

4.3. *The class-1894 named scenario*

The provincial design event cannot appear in the catalog: the rank-sampling ceiling and the fitted marginal jointly place an 1894-class discharge outside the sampled space. It is therefore priced as a named scenario outside the probabilistic framework: the official design discharge through the standard forcing machinery, its return period cited from the official study (NHC, 2008b) as a scenario attribute rather than as a coordinate of the exceedance curves. The scenario supplies what the catalog structurally cannot: the monetary consequence of the event the province already regulates against, priced by the same exposure, vulnerability, and aggregation machinery as everything else in this paper.

The forcing imposes a quasi-steady ramp confined to the spinup followed by a plateau of 18,900 m³/s at Mission, matching the provincial constant-discharge convention; the coastal boundary carries the astronomical high spring tide of the 70-year harmonic reconstruction (5.039 m Chart Datum, consistent with the catalog’s harmonic peak of 5.05 m) with zero net tidal residual; three variants span SLR 0, +0.77, and +1.50 m, simulated directly with the production hydrodynamic configuration.

The inventory response to the design event is: 5 assets materially wet (excess > 0.01 m) of 361 in scope at SLR 0, 14 at +0.77 m, and 47 at +1.50 m. The maximum excess at present sea level is 2.47 m at Barnston Island Pump Station, a station of an island dyking district whose exposure to the design event is expected; at +1.50 m the maximum is 4.06 m at Annacis Island Wastewater Treatment Plant, whose progression from zero excess at SLR 0 through 0.16 m at +0.77 m to 4.06 m at +1.50 m is the scenario-side confirmation of the threshold of Section 4.2.

The design-profile anchor check gives, at SLR 0 and in the sense model minus design: Steveston +0.90, New Westminster −1.00, Port Mann −1.61, Mission +0.49 m. The model brackets the official profile rather than sitting uniformly below it: the free interior underpredicts, consistent with the documented conservative bias of the hydrodynamic chain, while Steveston reflects a deliberately more severe coastal convention than the profile’s. The finding that the bias direction is not spatially uniform is carried into Section 5; the verification detail

behind the anchors (the channel-centre transect at Mission and the upstream decay of the sea-level transfer across variants) is provided in the Supplementary Information.

Priced through the production chain, the scenario yields CAD 4.8 M at present sea level, CAD 193.5 M at +0.77 m, and CAD 2,030 M at +1.50 m (central values). Three readings follow. First, for the protected critical infrastructure of this inventory, and conditional on the static crest representation, the provincial design event is today a minor event: a century of dike construction, initiated by this very flood, is visible in a single number. Second, the same hydraulic event scales by a factor of approximately 40 under +0.77 m and approximately 420 under +1.50 m of relative sea-level rise: nothing about the flood changes, only the freeboard that receives it. Third, the scenario losses are drawn as horizontal lines on the tier exceedance curves, letting the reader place the official design event against each probabilistic curve directly while the two frequency frameworks remain, deliberately, unmerged.

4.4. *Decomposing compound tail risk: statistical dependence versus hydrodynamic coupling*

This section reports the decomposition of the compound contribution to tail risk into its statistical-dependence and hydrodynamic-coupling components, under the design of Section 3.8. The impact calculations are run under both the production compound event sets and the counterfactual sets (the permutation catalogs of Experiment 1 and the four-leg scenario set of Experiment 2), and the aggregate loss exceedance curves and paired event losses are compared directly. The question the section answers is what fraction of aggregate tail losses exists because the drivers are statistically dependent, and what fraction because the estuary couples them non-linearly, by tier.

Replacing the fitted copulas with an independence counterfactual that preserves the empirical margins exactly (block permutation at the base-event level; three seeds) changes tail losses ($T > 100$ yr) by between 0.1 and 1.7 percentage points depending on tier and permutation draw, with no tier showing a contribution robustly distinguishable from zero across seeds. The estimated contribution of statistical dependence between river discharge and storm surge is therefore small and not robustly distinguishable from zero across permutation seeds, with the upper confidence bound remaining at approximately 3 percent of tail losses under this experiment. Average annual losses move by less than 1 percent under the same counterfactual, as expected when margins are preserved, and the headline-tier PML-100 moves from its production value (Table 2) to 1,856 CAD M, a shift of 0.1 percent. This result is additionally robust by construction to the frequency information-set gap declared as Mechanism 1 in Section 5: because the permutation preserves the exact empirical marginal draws, any recalibration of the fitted marginals displaces the production catalog and its independence counterfactual identically, leaving the contrast between them unchanged to first order, and a historically thickened marginal tail would raise single-driver severities and thereby operate in the direction of the marginal-dominance conclusion rather than against it.

The four-leg differences-in-differences design over the twenty most extreme simulated events isolates the net hydrodynamic interaction from the tier-and-tide background loss, evaluated with both drivers at their seasonal medians under the event's own tidal instant and sea-level offset. The interaction is negligible for the median event (+0.1 percent of event loss) but seasonally asymmetric: indistinguishable from zero in freshet-driven events (−0.01 percent, 95 percent bootstrap interval −0.19 to +0.14), and positive in surge-driven events (mean +13.7 percent, bootstrap interval +0.3 to +29.1, reaching approximately 30 percent

in individual events); the pooled mean over the twenty events is +6.9 percent (bootstrap interval +0.1 to +15.6). The intervals describe resampling variability within the deliberately selected extreme-event set, not confidence intervals for a population parameter. A parallel decomposition in station water levels rather than losses shows that superposition is nearly linear for the median wetted station (median interaction +0.0 cm, interquartile range within ± 0.1 cm) while individual stations in surge events accumulate up to 1.3 m of interaction-driven water, so the loss-level asymmetry is locally hydrodynamic and globally threshold-mediated: small water-level interactions are amplified where they cross protection freeboards, and are immaterial elsewhere. In this terminology, the water-level decomposition measures hydrodynamic interaction proper, while the loss-level term of Equation (4) measures that interaction propagated through the non-linear damage and threshold chain. The background term itself accounts for a median of 90 percent of compound event losses across the three tiers: under relative sea-level rise, the tier-and-tide state with unexceptional drivers already produces most of the loss, a structural property of a protected system rather than an artifact of the decomposition.

Table 4: Decomposition of compound tail risk by tier. ΔL_{stat} is the population tail share ($T > 100$ yr) under the primary permutation seed, with paired-bootstrap 95% CIs; the range across three permutation seeds is 1.34, 0.16, and 0.59 percentage points at +0.77, +1.00, and +1.50 m respectively, and no tier is robustly distinguishable from zero across seeds. The event-level interaction row summarises the four-leg design ($n = 6, 8$, and 6 events per tier); the pooled estimate in the text is primary and the per-tier values are indicative only.

	SLR_B (+0.77 m)	SLR_A (+1.00 m)	SLR_C (+1.50 m)
ΔL_{stat} share of tail loss (%)	1.71 [0.49, 2.93]	0.24 [−0.08, 0.52]	1.02 [−0.10, 2.15]
PML-100, independence counterfactual (CAD M)	1,363.9	1,856.0	2,455.8
Mean event-level interaction I/L_{comp} (%)	+11.8 [0.1, 33.4]	+10.4 [−0.2, 22.9]	−2.9 [−9.4, 0.6]

The event-level quartets, the counterfactual exceedance curves for all three permutation seeds, the water-level decomposition, and the reconciliation of the three-leg subtraction against the four-leg design (the three-leg form double counts the background and misreports the interaction as a median −90 percent) are provided in the Supplementary Information.

The tier structure of the decomposition is read against the concentration results of Section 4.2: because the tier AAL is dominated by threshold behaviour at a small number of assets, the share attributable to statistical dependence at a given tier reflects how much of the joint tail sampling moves those specific assets across their thresholds, and the interaction term isolates what the estuary’s non-linear response adds beyond the linear superposition of the single-driver excesses over background. The population decomposition is reported for tail events only (return period above 100 years), on the rank clock with analytic labels per Equation (1).

Both results are conditional on the declared model chain. The interaction estimates inherit the characterized hydrodynamic level bias of Section 3.4: the four-leg design cancels systematic level error to first order within each quartet, since all four legs share one model, one geometry, and one bias, but the residual bias is threshold-mediated, and under-predicted water levels imply under-counted freeboard crossings, so its sign is known. The positive interaction shares reported above are accordingly lower bounds with respect to that bias, while the negligible interaction at the median event is conditional on it; the qualitative

seasonal asymmetry, in which elevated coastal water impedes discharge conveyance while the reverse coupling is weak, is mechanistic and expected to be robust to level correction even where the magnitudes are not. The chain further represents flood defences as static crests without geotechnical fragility. Failure mechanisms sensitive to loading duration or to joint loading histories, rather than to peak water level, could restore a larger role for driver dependence and for hydrodynamic interaction; quantifying them requires a fragility-enabled defence representation and is identified as future work.

5. Declared systematic bias and limitations

The physical chain of this assessment carries a documented, systematic, directional underestimation of losses, concentrated in the tail. Five mechanisms constitute it, each with a declared sign; no correction is applied to the production results, which therefore constitute a lower bound with respect to these five mechanisms. Directional bias is not total uncertainty: the footprint-sampling and valuation envelopes are two-sided, and are bracketed where they matter (Sections 3.2 and 4.2) rather than folded into the bound. This section states the mechanisms, the quantitative bounds available for each, and the remediation path, in that order.

Mechanism 1: frequency information-set gap. The fitted freshet marginal uses the instrumental Mission record only ($n = 53$, 1965–2024), which contains neither historical flood of record; the official curve, incorporating them through the Expected Moments Algorithm, sits higher by approximately 15% at $T = 100$ and approximately 22% at $T = 500$. Sign: losses of a given magnitude are assigned frequencies that are too low; tail losses are understated in frequency, not severity. The historical-information sensitivity (refitting the marginal with the Expected Moments Algorithm to bound the induced movement of the tier curves) is specified as future work.

Mechanism 2: static, overtopping-only flood defences. Dike crests are static geometry and the only represented failure mode is overtopping; structural failure (saturation, scour, seismic), operational failure (gates, pumps), and breach dynamics are excluded. Sign: one-sided; any structural failure mode moves losses, and the Annacis threshold of Section 4.2, earlier, never later.

Mechanism 3: rank-sampling ceiling. Generator uniforms are drawn on $(0.001, 0.999)$, so no marginal can be sampled beyond its 0.999 quantile; the catalog maximum freshet discharge (16,086 m³/s) is the empirical realisation of the ceiling, an analytic ~ 990 -year event, and 1894-class discharges are unreachable by construction. Sign: the extreme tail of the driver space is truncated; catalog-based tail losses are understated. The named scenario of Section 4.3 supplies the complementary, non-probabilistic answer for the design event itself.

Mechanism 4: absence of secondary uncertainty. Each impact function is a deterministic curve and the loss given an event is a point value, so the year-loss tables of this chain sit one compression level below the interior of a vendor catastrophe model: no conditional severity distribution per cell is carried, and the dispersion of severity is compressed, parameter ranges being exercised in sensitivity analysis rather than propagated as secondary uncertainty. Sign:

compressing the conditional severity dispersion thins the simulated loss distribution’s tail; tail metrics are understated.

Mechanism 5: spatially non-uniform hydrodynamic level bias.. The production hydrodynamic configuration under-predicts peak water levels at the tidal-dominated gauges by 0.74 to 1.01 m against the June 2012 event, with the under-prediction growing monotonically with water level, while the scenario anchor check shows the model bracketing the official design profile rather than sitting uniformly below it: the free interior under-predicts, boundary-adjacent locations may not. Sign: under-predicted levels propagate to under-predicted depths and losses over the interior where most of the exposed inventory sits; the correction specified at the loss-aggregation stage (per-gauge, chainage-interpolated, with null and uniform envelope variants) is a declared design whose production run is future work, and its spatial uniformity is itself a declared approximation. Because the correction is calibrated on a single moderate freshet and the misfit grows with water level, the specified correction under-states the correction appropriate at catastrophic stages, reinforcing rather than weakening the lower-bound interpretation.

Interpretation rule.. The pre-specified rule is asymmetric, mirroring the conservative direction of the bias: conclusions that hold under the uncorrected chain are robust by construction, since any correction can only raise losses; conclusions that would appear only under corrected variants are not asserted in this paper.

Jensen-inequality bracket.. Depth–damage curves are scalar-to-scalar, and polygon-based extraction of extended assets at a single point introduces a Jensen-inequality error of unknown sign in the absence of an intra-asset value distribution. The defensible treatment is a min–max sensitivity bracket over the asset footprint, and Annacis Island Wastewater Treatment Plant is its named subject: with 83.5% of the +1.50 m tier AAL on one polygon, the bracket over its footprint is the sensitivity of the aggregate at that tier. The bracket is executed in Section 4.2: it is wide at the AAL level and at the lower-tier tail, and immaterial at the deep tail of the highest tier, so this mechanism is quantified rather than merely declared, and it does not overturn the regime conclusion.

Omission census: Mission City Station.. One asset is excluded from the tier product because the cells that physically correspond to the asset lie in the river channel, structurally inadmissible in the reduced-order domain. The omission is bounded by census rather than assumption: the catalog maximum discharge (16,086 m³/s in every tier) sits below the 18,900 m³/s class required to wet the corrected asset, zero of 200,000 events per tier reach it, and the exclusion is doubly locked by the rank-sampling ceiling and the analytic rarity ($T \sim 10^6$ yr) of the required discharge. The asset’s hazard statement is carried by the named scenario (0.40, 0.46, and 0.54 m of excess across the SLR variants).

Sensitivity context.. The dominant sensitivity axis is known from the platform’s own uncertainty analysis: impact-function threshold shift carries a total-order Sobol’ index of approximately 0.95 for average annual impact (Kropf et al., 2022), so threshold-like parameters (equipment elevation, datum uncertainty, and the defence crests of Mechanism 2) are expected to dominate the uncertainty budget of any chain of this architecture. Both delta

subsidence and compound flooding act by lowering effective asset-level thresholds, which is why they are treated as core structure rather than sensitivity variants in this assessment.

Exposure valuation uncertainty.. Valuation uncertainty is declared through the documented per-asset envelope rather than through a uniform scaling: the census lineage carries CAD 29.33 B (low), 55.61 B (central), and 103.90 B (upper), and all loss metrics are computed under the three variants from one shared damage-ratio computation (Section 3.7). A uniform $\pm 30\%$ scaling of all replacement values would move AAL and PML by exactly $\pm 30\%$ by homogeneity of degree one, and is therefore not run as a separate experiment; the heterogeneous envelope is the informative statement. A small number of unit-cost cells rest on documented fallbacks where primary sources could not be located, and remain the lowest-confidence entries of the valuation; they are flagged in the Supplementary Information rather than silently accepted.

Remediation path.. The residual hydrodynamic misfit is hypothesised to originate substantially in channel geometry: the provincial hydraulic model (Qu, 2024) rests on 761 surveyed cross-sections whose source data are not publicly available, and a future rebuild incorporating that geometry, in collaboration with the data holders, is the designated remediation. The remediation is partial (it addresses the hydrodynamic mechanisms, not the information-set, ceiling, or secondary-uncertainty mechanisms, which require separate treatment) and non-blocking: the present results are published as the declared lower bound with respect to the five characterized directional mechanisms, and no conclusion of this paper is conditional on the remediation being carried out. The attainable accuracy on this reach is demonstrated rather than hypothetical: an independently developed SFINCS configuration for the same estuary, incorporating the surveyed channel-geometry lineage and calibrated against four events including a 2025–2026 field campaign, achieves peak water-level root-mean-square errors of 0.13–0.25 m (De Jong, 2026), although that accuracy reflects calibration and tributary resolution as well as channel geometry.

Scope limitations.. Regime C hybrid rain-on-snow events are out of catalog scope (Section 2), an exclusion empirically supported under the present climate by De Jong (2026), whose transitional-season scenario ensembles for this estuary proved less severe than their freshet- and storm-season counterparts, while the projected growth of hybrid events under warming remains outside that support; the inventory is a census within five categories defined ex ante, and reported metrics are a lower bound with respect to airport and other excluded sectors (Section 3.2); the recomputation of the aggregate curves with the sensitivity-scope public-asset category included is future work.

6. Discussion

Set against the Canadian lineage, the contribution of this assessment is the layer that lineage stops short of. Expert-judgment resilience characterization (Chang et al., 2014) and disruption response curves over deterministic events (Murdock et al., 2018) answer how infrastructure systems degrade and recover; a probabilistic monetary loss chain answers how often and how much, in units that owners, insurers, and regulators consume, and over a frequency range (100,000 simulated years per tier) that deterministic event sets cannot reach.

The comparison with the global products is not a competition of accuracy but of structure: a per-peril global layer cannot represent the joint tail of an estuary where discharge, residual, tide, and mean sea level interact, and the present chain exists precisely to resolve that structure, from calibrated dependence through explicit protection geometry to asset-level loss.

Three findings deserve emphasis beyond their numerical statement. First, the present-day AAL-to-PML-100 ratio of 0.018 reads as functioning protection, not as low risk: the same portfolio migrates to 0.702 under +1.50 m of relative sea-level rise, and the migration is driven by identifiable assets crossing identifiable freeboard thresholds rather than by a smooth inflation of the curve. Second, the concentration structure is dynamic: the attribution shifts between the two dominant wastewater plants as sea level rises, so a single-tier reading of concentration would mischaracterize the system. Third, the named 1894-class scenario prices the provincial reference event at a level that a century of dike construction has made minor today, and shows the same event scaling by two orders of magnitude under the delta-adjusted tier with nothing about the flood itself changing; the probabilistic and scenario products tell one story from two independent constructions.

For regional compound-flood modelling, the decomposition design carries a transferable point. Compound studies typically report a lumped compound effect, leaving open whether the tail is a property of the copula or of the estuary. Where a process-based hydrodynamic model is part of the chain, the two are separable by construction: an independence-copula counterfactual isolates what statistical dependence contributes, and single-driver scenarios isolate what non-linear coupling adds beyond superposition. In a system whose tier losses are governed by threshold behaviour at specific assets, the split matters practically, because dependence operates through the frequency with which joint forcing crosses those thresholds while coupling operates through the water surface itself; adaptation levers differ accordingly. The framework requires only that the surrogate or model be defined over physical predictors rather than copula ranks, a design property worth adopting in any chain that anticipates counterfactual analysis.

The per-tier, asset-resolved metrics of Section 4 are also the quantitative layer that physical-risk disclosure and adaptation-finance decisions presuppose, with the declared lower-bound orientation itself decision-relevant.

The boundaries with the companion analyses are deliberate: verification of the extraction chain, translation of the year-loss tables into credit-risk quantities, and sensitivity to the choice among damage-function families are each developed elsewhere, and none of those results is anticipated here.

Read together, the compound decomposition, the concentration structure, the 1894 scenario, and the tier migration of the risk metrics are manifestations of one mechanism rather than four findings: relative sea-level rise consumes freeboard, raising the background loss that any tidal realisation produces; marginal extremes then cross protection thresholds they previously cleared, which concentrates losses at the specific assets whose thresholds fall first; and localised hydrodynamic interactions, negligible on average, become material precisely where that crossing occurs. The 1894 design event states the same mechanism in a single number: the flood does not change; the freeboard that receives it does. The mechanism is not specific to the Fraser: any protected lowland whose defences were dimensioned against a historical water-level distribution undergoes the same regime transition as relative sea-level

rise consumes the margin between routine water and crest, and the metrics developed here (the AAL-to-PML ratio trajectory and asset-level threshold crossings) transfer directly as its diagnostics.

The relationship of this chain to CanFlood, the flood risk toolbox recommended in the federal flood damage estimation guidelines (NRCan and Public Safety Canada, 2021) and the natural reference point for Canadian practice, deserves an explicit statement. The two are complementary rather than competing: CanFlood is deliberately designed as a flexible risk estimation platform that receives hazard surfaces and vulnerability functions supplied by the modeller, and this chain produces precisely those inputs, so a CanFlood implementation of this assessment would consume the same SFINCS water surfaces and apply the same curve-times-value arithmetic that the present engine executes and verifies against an independent implementation to machine precision (Section 3.7). The layer where a substantive comparison is possible is the vulnerability functions, and that comparison has been carried out as a provenance and coverage census. It is reassuring in both directions: for the utility occupancies that dominate the portfolio AAL, the function collection distributed with CanFlood draws on the same HAZUS utility-facility tables from which the production library was compiled under the sealed provenance keys of Section 3.6, with the wastewater-treatment and both pumping-plant curves numerically identical at matched depths of 1 and 3 m and the air-insulated substation curve differing only through its documented shallow-depth adaptation, so the production library and the national reference share their quantitative foundation where it matters most. Where the production curves depart (gas-insulated substations, open potable-water treatment), the departures are engineering adaptations recorded with rationale in the compilation report; and for the port, transit, and traction-power occupancies, which the distributed collection, with its current building-stock focus, does not yet cover, the custom compilation documented in Section 3.6 is offered as a contribution toward that coverage. A full cross-tool run was therefore not pursued, not because the comparison lacks value but because, under shared hazard, shared arithmetic, and a shared curve foundation, it would primarily test implementation equivalence rather than provide an independent benchmark of the physical-risk chain; the census and its matched-depth table are provided in the Supplementary Information.

Several extensions are specified but not executed. The fiscal-gap analysis, overlaying the aggregate exceedance curves on documented response capacity, is designed and its inputs exist. The business-interruption design of Section 3.2 awaits execution with its network-based cascade validation. The historical-information sensitivity, refitting the freshet marginal with the Expected Moments Algorithm to bound the movement of the tier curves attributable to the information set, addresses the first bias mechanism directly. The aggregate curves are to be recomputed with the sensitivity-scope public-asset category included. A second-generation chain incorporating verified channel geometry is the designated partial remediation of the hydrodynamic mechanisms.

7. Conclusions

This paper set out to quantify the probabilistic compound coastal–fluvial flood risk of the critical infrastructure of the Fraser River delta, and to do so under an epistemology suited to a protected system: drivers sampled where protection cannot filter them, outcomes computed

by physics with the protection represented explicitly, and every systematic bias of the chain declared with its sign.

The sealed production results support five statements. First, the present-day system behaves as a protected system should: an average annual loss of CAD 0.374 million against CAD 55.6 billion of exposure, an AAL-to-PML-100 ratio of 0.018 below the conventional catastrophe band, and annual losses carried by rare overtopping events. Second, that state is not robust to relative sea-level rise: the average annual loss multiplies by roughly three orders of magnitude across the tier structure, the ratio migrates to 0.702, and nearly every simulated year carries positive loss at the upper tiers, with the migration localised asset by asset rather than distributed. Third, the losses concentrate: two wastewater treatment plants jointly carry 97.5% of the headline-tier average annual loss, the dominant asset crosses a hydraulically coherent freeboard threshold between the +1.00 and +1.50 m tiers under the production sampling convention, the footprint bracket shows the regime shift, though not the attribution detail, is robust to that convention, and the concentration structure itself shifts with sea level. Fourth, the provincial 1894-class design event, priced through the same machinery as a named scenario, is a minor event for this inventory today and scales by factors of approximately 40 and 420 under the +0.77 and +1.50 m tiers, confirming from an independent construction that the driver of regime change is freeboard consumption, not hazard change. Fifth, compound tail risk in this system is not primarily generated by statistical dependence between discharge and coastal extremes: removing that dependence changes tail losses by amounts not robustly distinguishable from zero, bounded above by approximately 3 percent. Net hydrodynamic interaction is likewise negligible for the median extreme event but becomes material in surge-driven events, reaching approximately 30 percent of event loss, and the background state alone accounts for a median of 90 percent of compound event losses under sea-level rise. Compound tail risk is therefore predominantly a marginal-severity and freeboard phenomenon, with localised process interaction amplified where protection thresholds are crossed.

These statements rest on a chain whose five characterized bias mechanisms share a conservative sign; the reported loss metrics are therefore lower bounds with respect to those mechanisms, while sampling and valuation uncertainty are bracketed separately; the interpretation rule is asymmetric by design, and no conclusion of this paper depends on a corrected variant. What the assessment offers Canadian practice is a template as much as a number: a census-scoped, driver-domain, process-based chain whose conventions, verification lineage, and declared limitations are documented at the resolution a reader needs to reuse, audit, or remediate it.

References

- Altus Group (2025). *Canadian Cost Guide (2023, 2024, 2025, 2026 editions)*. Altus Group Limited. Grey literature: publicly distributed industry guide (Vancouver column). <https://www.altusgroup.com/insights/canadian-cost-guide/>
- Aznar-Siguan, G. & Bresch, D.N. (2019). CLIMADA v1: A global weather and climate risk assessment platform. *Geoscientific Model Development*, 12(7), 3085–3097.
- Province of British Columbia (2016). *Flood Hazard Area Land Use Management Guidelines* (amended). Victoria, BC.

989 BC Hydro, British Columbia Utilities Commission (2025). *West End Substation Project —*
990 *Certificate of Public Convenience and Necessity (CPCN) Initial Application*. BCUC Or-
991 ders G-278-25 (2025), G-80-26 (2026). Grey literature: BC public regulatory record; CAD
992 1.028B, 6,406 m2 underground GIS. [https://www.ordersdecisions.bcuc.com/bcuc/](https://www.ordersdecisions.bcuc.com/bcuc/orders/en/523141/1/document.do)
993 [orders/en/523141/1/document.do](https://www.ordersdecisions.bcuc.com/bcuc/orders/en/523141/1/document.do)

994 Bates, P.D., Horritt, M.S., & Fewtrell, T.J. (2010). A simple inertial formulation of the
995 shallow water equations for efficient two-dimensional flood inundation modelling. *Journal*
996 *of Hydrology*, 387(1–2), 33–45.

997 Bevacqua, E., Maraun, D., Vousdoukas, M.I., Voukouvalas, E., Vrac, M., Mentaschi, L.,
998 & Widmann, M. (2019). Higher probability of compound flooding from precipitation
999 and storm surge in Europe under anthropogenic climate change. *Science Advances*, 5(9),
1000 eaaw5531.

1001 Brice, B., et al. (2021). Paleoflood evidence of late-Holocene Fraser River freshets exceeding
1002 the instrumental record. *Environmental Research Letters*.

1003 Bryant, S. & Davies, T. (2022). Stochastic object-based flood damage assessment (SOFDA):
1004 A novel Monte Carlo framework for flood risk modelling. *Journal of Flood Risk Manage-*
1005 *ment*, 15(4), e12826.

1006 Chang, S.E., McDaniels, T., Fox, J., Dhariwal, R. & Longstaff, H. (2014). Toward disaster-
1007 resilient cities: Characterizing resilience of infrastructure systems with expert judgments.
1008 *Risk Analysis*, 34(3), 416–434.

1009 Couasnon, A., Eilander, D., Muis, S., Veldkamp, T.I.E., Haigh, I.D., Wahl, T., Winsemius,
1010 H.C., & Ward, P.J. (2020). Measuring compound flood potential from river discharge and
1011 storm surge extremes at the global scale. *Natural Hazards and Earth System Sciences*, 20,
1012 489–504.

1013 Curry, C.L., Islam, S.U., Zwiers, F.W., & Déry, S.J. (2019). Atmospheric rivers increase
1014 future flood risk in Western Canada’s largest Pacific River. *Geophysical Research Letters*,
1015 46(3), 1651–1661.

1016 Dashtgard, S.E., et al. (2024). Vulnerability of the Fraser River delta plain, British Columbia,
1017 to relative sea-level rise. *Canadian Journal of Earth Sciences*, 61(5), 543–557.

1018 De Jong, T. (2026). Probabilistic assessment and seasonal characterization of compound flood
1019 hazards in the Fraser River estuary. MASc thesis, The University of British Columbia,
1020 Vancouver. <https://doi.org/10.14288/1.0455963>

1021 Eilander, D., Couasnon, A., Leijnse, T., Ikeuchi, H., Yamazaki, D., Muis, S., Dullaart, J.,
1022 Haag, A., Winsemius, H.C., & Ward, P.J. (2023a). A globally applicable framework for
1023 compound flood hazard modeling. *Natural Hazards and Earth System Sciences*, 23(2),
1024 823–846.

1025 Eilander, D., et al. (2023b). HydroMT-SFINCS: automated SFINCS model building. Deltares
1026 open-source software. https://github.com/Deltares/hydromt_sfincs.

1027 Fox-Kemper, B., et al. (2021). Ocean, cryosphere and sea level change. In *Climate Change*
1028 *2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assess-*
1029 *ment Report of the IPCC*. Cambridge University Press.

1030 Fraehr, N., Wang, Q.J., Wu, W., & Nathan, R. (2022). Upskilling low-fidelity hydrodynamic
1031 models of flood inundation through spatial analysis and Gaussian process learning. *Water*
1032 *Resources Research*, 58(8), e2022WR032248.

1033 Fraehr, N., Wang, Q.J., Wu, W., & Nathan, R. (2023). Development of a fast and accu-
1034 rate hybrid model for floodplain inundation simulations. *Water Resources Research*, 59(6),
1035 e2022WR033836.

1036 Fraehr, N., Wang, Q.J., Wu, W., & Nathan, R. (2024). Assessment of surrogate models
1037 for flood inundation: The physics-guided LSG model vs. state-of-the-art machine learning
1038 models. *Water Research*, 252, 121202.

1039 Genest, C., Rémillard, B., & Beaudoin, D. (2009). Goodness-of-fit tests for copulas: A review
1040 and a power study. *Insurance: Mathematics and Economics*, 44(2), 199–213.

1041 Gillett, N.P., et al. (2022). Human influence on the 2021 British Columbia floods. *Weather*
1042 *and Climate Extremes*, 36, 100441.

1043 Federal Emergency Management Agency (2025). *HAZUS-MH Multi-Hazard Loss Estimation*
1044 *Methodology, Version 7.0*. Washington, DC.

1045 Huizinga, J., de Moel, H., & Szewczyk, W. (2017). Global flood depth-damage functions:
1046 Methodology and the database with guidelines. *EUR 28552 EN*, Joint Research Centre,
1047 European Commission.

1048 Islam, S.U., Curry, C.L., Déry, S.J., & Zwiers, F.W. (2019). Quantifying projected changes in
1049 runoff variability and flow regimes of the Fraser River Basin, British Columbia. *Hydrology*
1050 *and Earth System Sciences*, 23(2), 811–828.

1051 Jalili Pirani, F. & Najafi, M.R. (2023). Nonstationary frequency analysis of compound flood-
1052 ing in Canada’s coastal zones using dynamic Bayesian copula model. *Coastal Engineering*,
1053 182, 104296.

1054 Kropf, C.M., Ciullo, A., Otth, L., Hallegatte, S., & Bresch, D.N. (2022). Uncertainty and sen-
1055 sitivity analysis for probabilistic weather and climate risk modelling: An implementation
1056 in CLIMADA v.3.1.0. *Geoscientific Model Development*, 15(18), 7177–7201.

1057 Lee, M. & Parfitt, B. (2022). *Reckoning: The 2021 Floods, Slides and Extreme Weather in*
1058 *BC—Costs and Consequences*.

1059 Leijnse, T., van Ormondt, M., Nederhoff, K., & van Dongeren, A. (2021). Modeling com-
1060 pound flooding in coastal systems using a computationally efficient reduced-physics solver:
1061 Including fluvial, pluvial, tidal, wind- and wave-driven processes. *Coastal Engineering*, 163,
1062 103796.

- 1063 Maduwantha, P., et al. (2026). Multi-day event windows for compound coastal–fluvial flood
1064 definition on multi-driver coasts.
- 1065 Martello, M.V., Whittle, A.J., & Lyons-Galante, H. (2023). Depth-damage curves for rail
1066 rapid transit infrastructure. *Journal of Flood Risk Management*, 16, e12856.
- 1067 Mazzotti, S., Lambert, A., van der Kooij, M., & Mainville, A. (2009). Impact of anthropogenic
1068 subsidence on relative sea-level rise in the Fraser River delta. *Geology*, 37(9), 771–774.
- 1069 Moftakhari, H.R., Salvadori, G., AghaKouchak, A., Sanders, B.F., & Matthew, R.A. (2017).
1070 Compounding effects of sea level rise and fluvial flooding. *Proceedings of the National
1071 Academy of Sciences*, 114(37), 9785–9790.
- 1072 Mühlhofer, E., Kropf, C.M., Aznar-Siguan, G., Bresch, D.N., & Riedel, L. (2024). Lines and
1073 polygons exposure handler in CLIMADA: Disaggregation and aggregation of impact data
1074 for spatially extended assets. *Environmental Research Communications*, 6, 015005.
- 1075 Mundie, J.H., & Bell-Irving, R. (1986). Predictability of the consequences of the Kemano
1076 hydroelectric proposal for natural salmon populations. *Canadian Water Resources Journal*,
1077 11(1), 14–25.
- 1078 Musselman, K.N., Lehner, F., Ikeda, K., Clark, M.P., Prein, A.F., Liu, C., Barlage, M.,
1079 & Rasmussen, R. (2018). Projected increases and shifts in rain-on-snow flood risk over
1080 western North America. *Nature Climate Change*, 8(9), 808–812.
- 1081 Murdock, H.J., de Bruijn, K.M. & Gersonius, B. (2018). Assessment of critical infrastructure
1082 resilience to flooding using a response curve approach. *Sustainability*, 10(10), 3470.
- 1083 Northwest Hydraulic Consultants (2008a). *Fraser River Hydraulic Model Update—Steveston
1084 to Hope*. Prepared for Fraser Basin Council, Vancouver, BC.
- 1085 Northwest Hydraulic Consultants (2008b). *Comprehensive Review of Fraser River Flood Hy-
1086 drology and Flows at Hope: Scoping Study*. Report prepared for the British Columbia
1087 Ministry of Environment.
- 1088 Northwest Hydraulic Consultants (2014). *Fraser River Design Flood Level Update: Hope to
1089 Mission*. Technical report prepared for the BC Ministry of Forests, Lands and Natural
1090 Resource Operations.
- 1091 Natural Resources Canada and Public Safety Canada (2021). Federal Flood Damage Esti-
1092 mation Guidelines for Buildings and Infrastructure, Version 1.0. Federal Flood Mapping
1093 Guidelines Series, General Information Product 124e. Government of Canada, Ottawa.
- 1094 Nirandjan, S., Koks, E.E., Ward, P.J., et al. (2024). A global open-access database of flood,
1095 earthquake, and tropical-cyclone damage and fragility functions for critical infrastructure
1096 systems. *Natural Hazards and Earth System Sciences*, 24(12), 4443–4470.
- 1097 Ohenhen, L.O., et al. (2026). Land subsidence along the coast of Canada from satellite InSAR
1098 and groundwater data. *Nature*, forthcoming.

1099 Pacific Northwest National Laboratory (2016). *Valuation of Electric Power System Services*
1100 *and Technologies* (PNNL-25633). Richland, WA.

1101 Pollack, A. (2023). *Flood Depth-Damage Functions from HAZUS and 2015 NACCS* [Data
1102 set]. Zenodo. DOI: 10.5281/zenodo.10027236.

1103 Pugh, D., & Woodworth, P. (2014). *Sea-Level Science: Understanding Tides, Surges,*
1104 *Tsunamis and Mean Sea-Level Changes*. Cambridge University Press.

1105 Qu, J. (2024). A Detailed 2D Hydraulic Model for the Lower Fraser River. In R. Gupta
1106 et al. (eds.), *Proceedings of the Canadian Society of Civil Engineering Annual Conference*
1107 *2022*, Lecture Notes in Civil Engineering 363, pp. 775–798. Springer. [https://doi.org/](https://doi.org/10.1007/978-3-031-34593-7_48)
1108 [10.1007/978-3-031-34593-7_48](https://doi.org/10.1007/978-3-031-34593-7_48)

1109 Ryan, C. & Bristow, D. (2026). Climate flood risk assessment for critical infrastructure using
1110 CLIMADA and Bayesian Networks: Application to Halifax, Nova Scotia. *Working paper /*
1111 *forthcoming*.

1112 Sharmin, N., Henstra, D., & Thistlethwaite, J. (2026). A systematic review of the impacts of
1113 climate change on critical infrastructure in Canada. *Environmental Management*, 76, 103.

1114 Wahl, T., Jain, S., Bender, J., Meyers, S.D., & Luther, M.E. (2015). Increasing risk of
1115 compound flooding from storm surge and rainfall for major US cities. *Nature Climate*
1116 *Change*, 5(12), 1093–1097.

1117 Williams, J.H., et al. (2025). Expert-derived flood damage curves for critical infrastructure
1118 network components. *Journal of Flood Risk Management*, 18, e70045.

1119 Williams, J.J., & Esteves, L.S. (2017). Guidance on setup, calibration, and validation of
1120 hydrodynamic, wave, and sediment models for shelf seas and estuaries. *Advances in Civil*
1121 *Engineering*, 2017, 5251902. <https://doi.org/10.1155/2017/5251902>

1122 Wing, O.E.J., et al. (2020). New insights into US flood vulnerability revealed from flood
1123 insurance big data. *Nature Communications*, 11, 1444.

1124 Wing, O.E.J., Bates, P.D., Quinn, N.D., Savage, J.T.S., Uhe, P.F., Cooper, A., Collings,
1125 T.P., Addor, N., Lord, N.S., Hatchard, S., Hoch, J.M., Bates, J., Probyn, I., Himsworth,
1126 S., Rodriguez Gonzalez, J., Brine, M.P., Wilkinson, H., Sampson, C.C., Smith, A.M., Neal,
1127 J.C. & Haigh, I.D. (2024). A 30 m global flood inundation model for any climate scenario.
1128 *Water Resources Research*, 60, e2023WR036460.

1129 Zscheischler, J., Seneviratne, S.I. (2017). Dependence of drivers affects risks associated with
1130 compound events. *Science Advances*, 3, e1700263. DOI 10.1126/sciadv.1700263.

1131 Zscheischler, J., et al. (2018). Future climate risk from compound events. *Nature Climate*
1132 *Change*, 8(6), 469–477.

1133 Zscheischler, J., et al. (2020). A typology of compound weather and climate events. *Nature*
1134 *Reviews Earth & Environment*, 1, 333–347.

Statements and Declarations

Funding. This work was supported by a Four Year Doctoral Fellowship from The University of British Columbia.

Competing interests. The author has no relevant financial or non-financial interests to disclose.

Data availability. The sealed risk metrics and reproduction materials will be deposited in a public repository upon resolution of source-data licensing for a derived elevation product; in the interim they are available from the author on reasonable request.

Author contributions. The author solely conceived and designed the study, performed all analyses, and wrote the manuscript.

Declaration of generative AI use. During the preparation of this work the author used Claude (Anthropic) to assist with wording review, LaTeX formatting, and code debugging. All scientific decisions, model designs, data analyses, and conclusions are the author’s own. The author reviewed and edited all content and takes full responsibility for the content of the published article.